An intrinsic treatment of ubiquity phenomena in higher-order computability models (Part I, draft version)

John Longley

April 19, 2017

Abstract

We introduce some simple conditions on typed partial combinatory algebras (viewed as models of higher-order computability) which suffice for an axiomatic development of some non-trivial computability theory, including various forms of the Kreisel-Lacombe-Shoenfield theorem, along with the author's earlier 'ubiquity' theorems which say that a wide range of computability models all give rise to the same 'hereditarily total' functionals. This results in a cleaner, more 'intrinsic' formulation of the latter theorems, which we are moreover able to generalize to cover certain computability models of a non-deterministic character. Within our axiomatic framework, we also investigate extensions of the ubiquity results to a much wider class of types, arising from simple types via *subset* and *quotient* constructions. This enables us to cover many types of interest that arise naturally in mathematical practice (e.g. in real or complex analysis), and establishes the existence of robust computability notions at such types.

In the present Part I, we introduce our axioms and the conceptual framework, and carry through the required proofs for first- and second-order types. The focus here is on the Kreisel-Lacombe-Shoenfield theorem, which we adapt to our axiomatic setting in various ways. In Part II, we will extend our results to type 3 and above, using the construction of Normann which we exploited in our earlier work on ubiquity.

1 Introduction and motivations

This paper makes some contributions to the general theory of typed partial combinatory algebras (TPCAs), viewed as models of higher-order computability as in [16, 14, 15]. In particular, we offer some new perspectives on the 'ubiquity' results of [14], where we showed that a wide range of TPCAs, which may differ considerably in the 'computable' partial operations of simple type that they support, nevertheless all agree on the class of hereditarily extensional total functionals over $\mathbb{N}$ that they give rise to. More specifically, it was shown in [14] that a large class of *full continuous* models all give rise to the Kleene-Kreisel Continuous Functionals, and a large class of *effective* models all give rise to the Hereditarily Effective Operations.

The contributions of the present paper (and its intended sequel) are of two kinds. Firstly, we identify a set of axioms on TPCAs which allows us to treat the abovementioned results (and others) as part of the 'intrinsic' theory of a given TPCA A, in contrast to the more 'extrinsic' approach taken in [14]. Secondly, we show that the ubiquity phenomenon — that is, the agreement of widely differing computability models on certain classes of total functionals — typically extends to a much larger repertoire of types, including many *subset* and *quotient* types such as appear frequently in ordinary mathematical practice. We discuss these two contributions in turn.

1.1 An axiomatic basis

As portrayed in [16], the theory of higher-order computability consists of some general theory applicable to TPCAs (or more specifically to λ-algebras), alongside more specialized results pertaining to particular models or computability notions. As a broad assessment of the present state of the subject, it would seem fair to say that the general theory mostly offers a organizing framework and a uniform treatment of some of the basic material, whereas most of the 'deep theorems' in the area concern the structure of individual models. A significant exception is provided by the ubiquity results of [14], which demonstrate the possibility of proving highly non-trivial theorems at a general level, thus dealing with a whole class of models in a single swoop.

In order to obtain non-trivial general results of this kind, it appears that one needs some stronger hypothesis than simply the axioms for a TPCA A endowed with a certain level of computational machinery. In [14], these additional hypotheses took the form of a requirement that A admitted a well-behaved *simulation* in some other well understood TPCA: specifically, we called A a *continuous* TPCA if it admitted a well-behaved simulation in Kleene's K_2, and an *effective* TPCA if it admitted a good simulation in K_1. It was shown that the existence of a suitable simulation was sufficient to allow certain well-known 'continuity' properties of K_2 or K_1 to be transferred to A, leading to the proofs of the ubiquity theorems.

Whilst the possibility of such general theorems was encouraging in principle, their formulation in terms of simulations was somewhat cumbersome and not entirely clean. Moreover, the proofs were somewhat complicated by the fact that rather than working simply with properties of a model A in itself, we were forced to work with a simulation in another model, so that two distinct computational 'tiers' had to be kept in play. (Actually, since the results in question concerned the simulation of yet another type structure within A, there were in practice three tiers involved.) It is therefore reasonable to ask whether a cleaner, more axiomatic approach is possible: can we identify some key properties of A *itself* which are implied by the existence of a certain kind of simulation, and which in turn suffice as a basis for the proofs of the ubiquity theorems?

One of the main contributions of the present work is the identification of a selection of relatively simple and clean axioms on a model A which, in various combinations, allow for the development of the main results of [14] (and further extensions of these), in conjunction with various forms of the *Kreisel-Lacombe-Shoenfield* theorem, as part of the 'intrinsic theory' of A (in other words, without reference to a simulation in some other model or any other structure external to A). Specifically, we shall introduce five main axioms, which we call Continuity, Enumeration, Normalization, Collection and Restriction (some of these come in more than one flavour).

Although the particular selection of axioms adopted here is to some extent engineered to support the specific proofs we have in mind, our longer term goal is the identification of a more general-purpose set of axioms that can serve as the basis for the clean mathematical development of a substantial body of general theory, redressing the imbalance we hinted at above. It is therefore encouraging that our axioms show signs of convergence with other conditions that have emerged as useful in our work elsewhere (see e.g. [16, Section 7.1]). There are also reminiscences of axioms previously considered, e.g. in synthetic domain theory [11] and synthetic computability theory [2], although the spirit of our present work is rather different (for instance, we do not here consider our datatypes as 'sets' within an appropriate topos).[1]

Whilst our objective is to render the use of simulations formally unnecessary, it is worth noting that some of our axioms can be seen as *inspired by* the typical scenario of a model A equipped with a simulation in K_2 or K_1. That is, we ask ourselves: if A *did* admit such a simulation, what traces of this fact would be discernible within A itself? This is perhaps a helpful light in which to view our axioms for Continuity and Enumeration, and even more so our Collection Axiom, a more novel condition which was specifically suggested to us by this way of thinking. (The relationships between our axioms and the more extrinsic, simulation-based view will be spelt out in Section 2.7.)

Having freed ourselves from the need to postulate a K_2- or K_1-simulation, our axiomatic development of the theory is in principle more general than before — although we do not see this as a major selling-point, since as far as we are aware, all particular models of interest that satisfy our new axioms also admit a simulation satisfying the conditions in [14]. In one respect, however, we make a genuine advance: we show how to generalize our treatment to embrace models of a 'non-deterministic' flavour, addressing a gap noted in [14]. (Indeed, the way forward here has perhaps become easier to see now that simulations has been removed from the picture.) The appropriate generalization, which involves non-deterministic variants of the Continuity and Enumeration Axioms, will be presented in Section 6.

One curious point that deserves mention is that not all of our axioms express properties that are stable under *simulation equivalence* of TPCAs (see [16, Section 3.3]) — in other words, they do not all manifest themselves as categorical properties of $\mathcal{M}od(A)$, the category of modest sets over A. On the one hand, one could argue that this is not a serious problem: since the *conclusions* we aim for are indeed properties of $\mathcal{M}od(A)$, all that really matters is that *some* model A' equivalent to A satisfies our axioms. On the other hand, one might see this non-robustness of our axioms as an indication that we have not yet put our finger on the essential properties that really explain the phenomena in question. Against this latter point, however, there is some experience to suggest that whilst $\mathcal{M}od(A)$ may be argued to embody what is essential about A from the perspective of *realizability*, a finer-grained view is desirable if we wish to consider A from the perspective of *computability* as we do here (see [5] and the discussion in [15]). Clearly, such questions deserve further consideration, but our present position is that whilst we are at the 'fieldwork stage' in gathering candidates for axioms and evidence for their value, we should not be constrained by the dictates of one particular categorical perspective on the nature of TPCAs.

In the present Part I, our emphasis is on setting up the framework and introducing the

[1]Our general enterprise of trying to develop portions of computability theory on an axiomatic footing may also recall earlier work in generalized or abstract computability theories, as in Kreisel [12], Moschovakis [17], Fenstad [8]. However, there appears to be little substantial connection between these earlier theories and our present endeavour, which is specifically addressed to a higher-order setting.

axioms, which we then use to prove ubiquity-style results at type level 2: this amounts to showing how various forms of the classical Kreisel-Lacombe-Shoenfield argument [13] can be carried over to our general setting. Mathematically this corresponds only to a tiny fragment of the results of [14]: the real substance of the ubiquity theorems, which depends on an ingenious construction due to Normann [18], only kicks in at type level 3. Nonetheless, the situation at type 2 presents a self-contained and already quite rich story that illustrates the potential of our axiomatic framework. It is also interesting to note that this portion of the theory requires only relatively weak computational machinery: at type 2 it is enough that A supports *iteration*, whereas at higher types a more general form of *recursion* is necessary. The task of carrying through the proofs for type 3 and above will be left for Part II.

1.2 Subset and quotient types

We now turn our attention to the ways in which our results go beyond those of [14].

We may begin by noting the fundamental role of Church's simple theory of types [4], not only as providing the basic infrastructure for our theory of TPCAs and indeed for many simply typed programming languages, but also as a possible setting for the foundation of mathematics. It is well known, for example, that much of ordinary number theory can be developed within *first-order* arithmetic (with just a single type for the natural numbers), whilst much of the theory of real and complex numbers can be formulated within second-order arithmetic (where one also has a type for functions on the natural numbers). Extending this to all finite type levels, we obtain *higher-order* arithmetic (typified by the systems PA^ω and HA^ω); and it is has frequently been asserted that this system, and even the first three or four levels of it, is sufficient as a foundation for the bulk of traditional mathematics. Indeed, a foundational framework of just this kind is employed in the HOL theorem proving system and its relatives [10]; in a somewhat more constructive context, the idea that a finite type structure over $\mathbb{N}$ provides a sufficiently rich universe for much of mathematical practice is also a theme of Feferman's work (e.g. [7]).

In practice, however, a typical mathematical development will use not just the plain simple types over $\mathbb{N}$, but a variety of *subset* and *quotient* types derived from these: the classical constructions of the rationals and the Cauchy or Dedekind reals provide a case in point. Indeed, derived types of this kind are a staple of the HOL system, where a mechanism for forming subset types is provided as a primitive, and a quotient type may be defined as a certain subtype of a powerset type. In the context of HOL, of course, this is really a matter of practical convenience rather than strict logical necessity: for instance, any function whose domain is a subset type $\sigma \subseteq \tau$ could in principle be adequately represented by a function on τ, since in classical logic there is no problem in extending a function on a subset of τ to one on the whole of τ. However, the situation is quite different if our intention is to model some constructive or semi-constructive system via some more restricted universe of *continuous* or even *computable* objects. In some situations, the existence of suitable extensions may be guaranteed by some explicit extension theorem or injectivity property, but there are also many naturally arising situations in which no such extension is available. We mention here a small selection of 'mainstream' examples:

1. The *minimization* or μ *operator*, defined on functions $f : \mathbb{N} \to \mathbb{N}$ with $f \neq \Lambda n.0$ and

yielding the smallest n for which $f(n) \neq 0$, is (Kleene) computable, but does not extend to a computable or even a continuous function on the whole of $\mathbb{N} \to \mathbb{N}$.

2. The computable function $x \mapsto 1/x : \mathbb{R}-\{0\} \to \mathbb{R}$ does not extend computably or even continuously to the whole of $\mathbb{R}$.

3. Let L be the space of closed curves in the plane (i.e. continuous functions from the unit circle S^1 to the plane $\mathbb{R}^2$), and let L' be the set of all $\gamma \in L$ that avoid the point $o = (0,0)$ (i.e. such that $o \notin \mathrm{Im}\ \gamma$). Then there is a well-defined function $w : L' \to \mathbb{N}$ which returns the *winding number* of a given curve γ about o, and this function is even computable in an appropriate sense. However, w does not extend to a continuous function $L \to \mathbb{N}$ for any reasonable topology on L, since we may smoothly interpolate between a curve with winding number 0 and one with winding number 1. Indeed, since the transition between 0 and 1 will typically happen abruptly as the curve passes over the point o, we are unable to give a continuous extension of w even if we replace the codomain $\mathbb{N}$ by $\mathbb{R}$.

   Similar situations will arise in connection with numerous other topological invariants.

4. Let D denote the closed unit disc in the complex plane, and δ its boundary. Let F denote the set of *real-differentiable* functions $D \to \mathbb{C}$ that are non-zero on γ, and let F' denote the set of *analytic* functions in F. Then there is a well-defined function $\sharp : F' \to \mathbb{N}$ where $\sharp(f)$ is the number of zeros of f within D (counting zeros by their multiplicity). Indeed, a standard result of complex analysis says that

$$\sharp(f) \;=\; \frac{1}{2\pi i} \int_\delta \frac{f'}{f} \;,$$

   which in turn shows that $\sharp$ is computable in an appropriate sense. However, an interpolation argument again shows that $\sharp$ cannot be continuously extended to a function $F \to \mathbb{N}$. (In this case, we *can* obtain a continuous extension if we replace the codomain by $\mathbb{R}$ — indeed, the integral formula above gives one way to do this.)

There are also many examples arising from computability theory itself: for instance, if $H \subseteq \mathbb{N}$ is the halting set, the 'halting time' function $H \to \mathbb{N}$ has no computable extension to the whole of $\mathbb{N}$.

The essential point here is that if we are given some universe of 'continuous' or 'computable' functionals only at the simple types over $\mathbb{N}$, then there is no way, without further conditions, to extrapolate from this to determine what the corresponding functions on all subsets of interest ought to be. There is therefore the danger that we may miss out on some important functions if we commit ourselves to working with total functionals of simple type.

We may now make the connection with our earlier discussion of TPCAs and ubiquity theorems. Regarding a TPCA A as providing some class of 'computable' entities of higher type, the total type structure over $\mathbb{N}$ obtained from A (equivalently the type structure over the evident object N within $\mathcal{M}od(A)$) can be seen as the class of 'A-computable' functionals of simple type, and the force of the ubiquity phenomenon is that this class is remarkably robust with respect to the choice of A. However, in the light of the above, we now see that this total type structure might fail to capture important aspects of what is 'A-computable' in the realm of subset and quotient types. It is therefore natural to ask what may be said about the A-computable functionals at such types (which are, after all, present as objects of $\mathcal{M}od(A)$), and in particular whether a similar ubiquity phenomenon holds for all or some of these new types.

Our present work breaks some new ground in this area. At the heart of our investigation is a concept of the *spectrum* of a TPCA A with numerals. This is in effect just the infinite tree representing the generation of a class of modest sets (starting from N) via just two operations: forming exponentials $- \Rightarrow N$, and taking arbitrary regular subobjects. The first of these, by itself, would lead us to the familiar class of *pure* types, and it is well known that these are sufficient to encode the whole universe of simple types (see e.g. [16, Section 4.2]). In a somewhat similar way, we will see in Section 7 that the spectrum of A determines the entire full subcategory of $\mathcal{M}od(A)$ generated from N by closing under finite products, exponentials, regular subobjects and quotients (we shall call this category the *envelope* of the spectrum). The idea here is to break free from the shackles of simple type structures and instead to work with a richer repertoire of types better suited to the demands of mathematical practice.

By studying various parts of the spectrum of A, we are able to learn something about the corresponding portions of the envelope that they give rise to. Specifically, we will show under various combinations of axioms on A that some substantial and interesting parts of the spectrum of A are *regular*, meaning in effect that they are determined purely by the set Δ of functions $\mathbb{N} \to \mathbb{N}$ present in A. (As we shall see, exactly *which* parts of the spectrum can be proved regular will depend on the particular combination of axioms adopted.) We thus obtain some robust classes of 'Δ-computable' functions for these parts of the spectrum, and hence for the portions of the envelope that they encode.

This in turn promises to yield robust computability notions for many types arising in mathematical practice, e.g. for spaces of continuous or analytic functions on a specified domain, or for spaces of operators acting on such functions. As mentioned above, in the present Part I we have only carried through this programme up to second-order types, and at this level the fruits of our labours are relatively modest (see Section 7.1). However, we fully expect that once our results have been extended in Part II to higher types, a wealth of interesting and non-trivial applications of this kind will be forthcoming, and the present paper paves the way for this.

1.3 Content and outline of paper

In Section 2 we present our axiomatic setup, establishing the basic notion of a *TPCA with iteration*, introducing in turn our axioms for Continuity, Enumeration, Normalization, Collection and Restriction (and their variants), and indicating the range of models that these axioms cover. We also show which of these axioms hold as consequences of the existence of a well-behaved simulation in K_2 or K_1. In Section 3, as a gentle warm-up, we show from the Enumeration Axiom alone that the set Δ of A-representable functions $\mathbb{N} \to \mathbb{N}$ completely determines the *1-spectrum* of A, that is, the types of level 1 within the spectrum of A. The core of the paper is in Sections 4 and 5, where we investigate (under various axiom combinations) how much of the *2-spectrum* is likewise determined by Δ. Section 4 considers this problem for types of the form $R \Rightarrow N$ where R is a subset of $N \Rightarrow N$. The key mathematical ingredient here is the original proof of the Kreisel-Lacombe-Shoenfield theorem ([13]; also [16, Section 9.2]) which we are able to adapt to our axiomatic setting in two slightly different ways. Section 5 extends the investigation to types $R \Rightarrow N$ where $R \subseteq Q \to N$ and $Q \subseteq N$; here we present some positive results and some counterexamples.

In Section 6, we show that a small adjustment to our framework allows to embrace many models of a *non-deterministic* flavour, and that a version of the KLS argument still goes through in this setting. As mentioned above, this represents a genuine advance over [14] where we were unable to treat such models. In Section 7, we introduce a general concept of *abstract spectrum*, and show that any such spectrum gives rise to a concrete cartesian closed category endowed with all possible 'subset' and 'quotient' types; in the concrete case of the spectrum of a model A this envelope coincides with a full subcategory of $\mathcal{M}od(A)$. We also give a simple illustration of the intended use of this concept, addressing the concept of computability for functions on subsets of $\mathbb{R}^n$. Finally, in Section 8, we mention some outstanding questions and the prospects for further work, including the extensions to higher types which we will undertake in Part II.

2 The axiomatic setup

We here introduce the structures we will be dealing with, along with the axioms which, in various combinations, will allow us to derive the results of this paper.

In Section 2.1 we recall some basic material on TPCAs from [16, Chapter 3], and establish a notion of 'TPCA with iteration' as a suitable baseline for our investigations. In Sections 2.2 to 2.6 we introduce, in turn, our axioms for Continuity, Enumeration, Normalization, Collection and Restriction (and their variants), informally discussing the role of each axiom within our theory and the range of models to which it applies. In Section 2.7 we consider which of our axioms follow from the existence of a certain kind of simulation in K_2 or K_1; this clarifies the relationship between our present work and the 'extrinsic' approach of [14].

As in [16], we adopt the following notational conventions in connection with potentially non-denoting expressions: $=$ denotes strict equality (the values of both sides are defined and they are equal); $\simeq$ denotes Kleene equality (if either side has a defined value then so does the other and they are equal); and $\downarrow, \uparrow$ denote definedness and undefinedness respectively.

2.1 TPCAs, numerals and iteration

We first define the class of 'models of generalized computation' that we will be working with. We refer the reader to [16, Chapter 3] for a more detailed exposition of the basic theory and for numerous examples.

Throughout this paper, we shall work with the set of *simple types* σ generated syntactically by the grammar

$$\sigma \ ::= \ \mathtt{N} \ | \ \sigma \to \sigma \ ,$$

where $\mathtt{N}$ is the type we shall use for representing the natural numbers. (As usual, we treat $\to$ as right-associative.) We write $\overline{k}$ for the *pure type* of level k: $\overline{0} = \mathtt{N}$, and $\overline{k+1} = \overline{k} \to \mathtt{N}$. We also write $\sigma^r \to \tau$ for the type defined by $\sigma^0 \to \tau = \tau$ and $\sigma^{r+1} \to \tau = \sigma \to (\sigma^r \to \tau)$; this should cause no confusion as we are not including product types in our system.

The following specializes Definition 3.1.16 of [16] to this setting:

Definition 1 *(i) A* partial applicative structure *A consists of an inhabited set* $A(\sigma)$ *for each type* σ*, and for each pair of types* σ, τ *a partial function* $\cdot_{\sigma\to\tau} : A(\sigma \to \tau) \times A(\sigma) \rightharpoonup A(\tau)$ *(called* application*, and treated as left-associative).*

(ii) A partial applicative structure A is a typed partial combinatory algebra*, or* TPCA*, if for all types* σ, τ, υ *there exist elements*

$$k \in A(\sigma \to \tau \to \sigma)\ , \qquad s \in A((\sigma \to \tau \to \upsilon) \to (\sigma \to \tau) \to \sigma \to \upsilon)$$

such that the following hold for all a, b, c, f, g *of suitable types:*

$$k \cdot a \cdot b \;=\; a\ , \qquad s \cdot f \cdot g\downarrow, \qquad s \cdot f \cdot g \cdot a \;\simeq (f \cdot a) \cdot (g \cdot a)\ .$$

The motivation for part (ii) of the above definition is that it is equivalent to a certain *combinatory completeness* property for partial applicative structures, which informally says that any function representable by means of a formal expression in the language of the applicative structure is also representable by an element within the structure itself. To make this precise, we consider *formal (applicative) expressions* over A, built up syntactically via formal application from typed *variables* x^σ and *constants* a^τ for each $a \in A(\tau)$, subject to the obvious typing constraints. We shall represent application within formal expressions by simple (left-associative) juxtaposition of sub-expressions, in contrast to semantic application for actual elements of A which will be denoted by '$\cdot$'. Given any formal expression e of type σ and a *valuation* ν assigning to each variable appearing in e an element of A of the appropriate type, we write $[e]_\nu \in A(\sigma)$ for the *value* of e with respect to ν defined in the obvious way; note that $[e]_\nu$ will be undefined if some of the applications prescribed by e are undefined in A with respect to ν.

In its general form, combinatory completeness may be stated as follows: for any formal expression e whose variables are among $x, y_0, \ldots, y_{r-1}$, we may construct a certain formal expression d with variables among $y_0, \ldots, y_{r-1}$ such that for any valuation ν for $y_0, \ldots, y_{r-1}$ and any value a for x, we have $[d]_\nu\downarrow$ and $[d]_\nu \cdot a \simeq [e]_{\nu, x\mapsto a}$. We denote this expression d using the meta-notation $(\lambda^* x.e)$ in order to indicate its significance: note that $[(\lambda^* x.e)a]_\nu \simeq [e[x \mapsto a]]_\nu$ for any ν. We shall in practice assume a certain fluency in the use of this principle for 'programming in TPCAs'.

We are interested in TPCAs that admit a good representation of natural numbers. For the present paper, we shall only require quite weak properties of our representation:

Definition 2 *A* TPCA with weak numerals *is a TPCA A equipped with a choice of distinct elements* $\widehat{0}, \widehat{1}, \widehat{2}, \ldots \in A(\mathtt{N})$ *such that the following hold:*

- *There exists* $suc \in A(\mathtt{N} \to \mathtt{N})$ *such that* $suc \cdot \widehat{n} = \widehat{n+1}$ *for all* $n \in \mathbb{N}$*.*
- *There exists* $rec \in A(\mathtt{N} \to (\mathtt{N}^2 \to \mathtt{N}) \to \mathtt{N} \to \mathtt{N})$ *such that for* $a \in A(\mathtt{N})$*,* $f \in A(\mathtt{N}^2 \to \mathtt{N})$ *and* $n \in \mathbb{N}$ *we have*

$$rec \cdot a \cdot f \cdot \widehat{0} \;=\; a\ , \qquad rec \cdot a \cdot f \cdot \widehat{n+1} \;\simeq\; f \cdot \widehat{n} \cdot (rec \cdot a \cdot f \cdot \widehat{n})\ .$$

For our present purposes, we take the choice of numerals to be part of the data for a model, and write $N \subseteq A(\mathtt{N})$ for the set of numerals.

Although $A(\mathtt{N})$ may contain other elements besides numerals, we shall typically be interested in whether or not a computation of type N yields a numerical result. Accordingly, we supplement our existing notation and terminology as follows. If e is an expression of type N, we shall say that e *converges* (and write $e \Downarrow$) if the value of e is defined as in within N; likewise, e *diverges* ($e \Uparrow$) if the value of e is either undefined or outside N. We will also write $e \cong e'$ to mean that e, e' either both diverge or converge to the same numeral. (Note that this is quite different in spirit from the treatment of convergence and divergence via *dominances* as in synthetic domain theory [11].)

As we shall see in Proposition 5, all our models of interest will admit diverging computations. Note, however, that there are many natural models in which evaluation of expressions is total (and diverging expressions will typically take a special value $\perp$), and many others in which N is the whole of $A(\overline{\mathtt{N}})$ (so that diverging expressions will never have a defined value). Our setup is designed to handle both of these situations (and others besides) in a uniform way.

In contrast to the usual definition of 'TPCA with numerals' (see [16, Section 3.3.4]) we do not require a recursor $rec_\sigma \in A(\sigma \to (\mathtt{N} \to \sigma \to \sigma) \to \mathtt{N} \to \sigma)$ for every type σ, but only for $\sigma = \mathtt{N}$. Although most naturally arising TPCAs with weak numerals also have numerals in the stronger sense, an important example of one that does not is the TPCA $\mathsf{SP}^{0,\mathrm{lbd}}$ of *left-bounded* closed sequential procedures, which can be seen as embodying Kleene's concept of *μ-recursiveness* (see [16, Section 6.3.3]).

We assume from now on that A is a TPCA with weak numerals, and allow ourselves to write $x : \sigma$ in place of $x \in A(\sigma)$. Clearly, every primitive recursive function $\mathbb{N}^r \to \mathbb{N}$ is representable by an element of A. We now take $\langle \cdots \rangle : \mathbb{N}^* \to \mathbb{N}$ to be some standard primitive recursive encoding of finite sequences of natural numbers, and shall make incessant use of the fact that the length, cons, head and tail operations for this encoding are A-representable. We take $(-)_i$ to be an operation that extracts the element at position i from a coded list when this exists, primitive recursively uniformly in i: thus $(\langle x_0, \ldots, x_{l-1} \rangle)_i = x_i$ when $i < l$.

We also note that A contains an element $ifeq : \mathtt{N}^4 \to \mathtt{N}$ such that for all $m, n \in \mathbb{N}$ and $a, b \in A(\mathtt{N})$ we have

$$ifeq \cdot \widehat{m} \cdot \widehat{n} \cdot a \cdot b \ = \ \begin{cases} a & \text{if } m = n \\ b & \text{if } m \neq n \ . \end{cases}$$

However, a corresponding operator $ifeq_\sigma : \mathtt{N}^2 \to \sigma^2 \to \sigma$ is not available for arbitrary σ.

Finally, we may encode the disjoint sum $\mathbb{N}+\mathbb{N}$ as $\mathbb{N}$ using e.g. the left and right injections $inl(n) = 2n$, $inr(n) = 2n+1$, noting that these are representable in A, and that there is an operator $case : \overline{1} \to \overline{1} \to \overline{1}$ satisfying

$$case \cdot f \cdot g \cdot \widehat{inl(n)} \ \simeq \ f \cdot \widehat{n} \ , \qquad case \cdot f \cdot g \cdot \widehat{inr(n)} \ \simeq \ g \cdot \widehat{n} \ .$$

We make use of this encoding in the next definition:

Definition 3 *Let A be a TPCA with weak numerals. We say A* has (ground-type) iteration *if it contains an element $iter : \overline{1} \to \overline{1}$ such that:*

$$\begin{aligned} iter \cdot f \cdot \widehat{n} \ &= \ \widehat{m} && \text{if } f \cdot \widehat{n} = \widehat{inl(m)} \\ iter \cdot f \cdot \widehat{n} \ &\simeq \ iter \cdot f \cdot \widehat{m} && \text{if } f \cdot \widehat{n} = \widehat{inr(m)} \ . \end{aligned}$$

The operation *iter* in effect gives us the power of while-loops manipulating ground type data. It is an easy exercise in programming in TPCAs to show that the existence of such an element *iter* is equivalent to the existence of a *minimization* or μ operator $min : \overline{2}$ such that whenever $g : \overline{1}$ and m satisfy $g \cdot \widehat{n} = \widehat{0}$ for all $n < r$ and $g \cdot \widehat{r} = \widehat{m+1}$, we have $min \cdot g = \widehat{r}$.[2] Although there has been a tradition in computability theory of using minimization as a primitive, we consider the notion of iteration to be more fundamental, and we are also eager to express our required levels of computational power in terms of familiar programming language constructs. It is a pleasing fact, for instance, that ground-type iteration will suffice for our analysis at type level 2, whereas at higher types a more powerful *recursion* operator will be required. The model $\mathsf{SP}^{0,\mathrm{lbd}}$ mentioned above is an example of a TPCA with ground-type iteration but without recursion even at the lowest useful level.

Relative to a choice of model A, we shall write N^N for the set of $g \in A(\overline{1})$ such that $g \cdot a \in N$ for every $a \in N$, and Δ_A, or just Δ, for the set of mathematical functions $\mathbb{N} \to \mathbb{N}$ representable by elements of N^N.

If A has iteration, it is easy to see that Δ_A contains every computable function $\mathbb{N} \to \mathbb{N}$, and indeed is *closed under Turing computation*: if $f_0, \ldots, f_{r-1} \in \Delta_A$ and $g : \mathbb{N} \to \mathbb{N}$ is Turing computable relative to $f_0, \ldots, f_{r-1}$ then also $g \in \Delta_A$. The set Δ_A will play a key role as an index of the computational power available in A: indeed, we shall see that under certain conditions, the class of A-representable functions at many other types is completely determined by Δ_A.

The class of TPCAs with weak numerals and iteration is quite wide: it embraces term models for a wide range of simply-typed programming languages (such as Plotkin's PCF and extensions thereof), complete partial order and lattice models such as the Scott, stable and strongly stable models, many 'intensional' models such as those arising in game semantics [1], and many 'relational' models such as the multiset model of [3]. Moreover, several of these models are often considered in both a 'full' flavour (with $\Delta = \mathbb{N}^{\mathbb{N}}$) and an 'effective' one (with Δ consisting of just the computable functions on $\mathbb{N}$). However, one may also construct models in which Δ is (for instance) the set of hyperarithmetic functions, or those computable in some specific oracle $h : \mathbb{N} \to \mathbb{N}$.

Furthermore, our framework also embraces all *untyped PCAs*. An untyped PCA is just a single set U equipped with a partial application $\cdot : U \times U \rightharpoonup U$, such that there exist $k, s \in U$ satisfying

$$k \cdot a \cdot b \;=\; a\,, \qquad s \cdot a \cdot b \downarrow\,, \qquad s \cdot a \cdot b \cdot c \;\simeq\; (a \cdot c) \cdot (b \cdot c)$$

for all $a, b, c \in U$. We may turn any such U into a TPCA A in our present sense simply by taking $A_\sigma = U$ for every σ, and $\cdot_{\sigma\tau} = \cdot$ for every σ, τ. It is well known that any untyped PCA automatically possesses numerals, iteration and indeed much more (see [16, Section 3.3]). Important examples of untyped PCAs include: Kleene's first and second models K_1 and K_2; the Scott graph model $\mathcal{P}\omega$ and similar relational models; the van Oosten model B for untyped sequential algorithms and similar untyped game models; and term models for various untyped λ-calculi.

[2] This is slightly stronger than the definition of minimization in [16, Section 3.3.5], because we do not require that $g \cdot \widehat{n} \Downarrow$ when $n \geq r$. We are also here adopting the convention that *min* searches for the first non-zero value rather than the first zero one, though this is inessential.

For the remainder of the paper, a *standard model with iteration* will mean a TPCA with weak numerals and ground-type iteration. This will provide our baseline notion of 'model of computation' on which we shall build by adopting several more specialized axioms. We now consider these additional axioms in turn, roughly in order of their significance.

2.2 Continuity

Our first axiom — and the most fundamental to our approach — will capture the idea that within a world of 'finitary' computation it is impossible to distinguish the function $\Lambda j.0$ from all other functions. One possible intuition for this, applicable to many models, is that any test that succeeds for $f = \Lambda j.0$ can only 'look at' the value of f on finitely many j, and must therefore also succeed on other functions that take the value 0 on these j. Another possible intuition, more appropriate in certain situations, is that being able to distinguish an infinite sequence of 0s from a sequence that at some point yields 1 would be tantamount to solving the halting problem.

We can capture this idea as follows. Let $\bot, \top \subseteq A(\overline{1})$ denote the sets

$$\begin{array}{rcl} \bot & = & \{\alpha \mid \forall j \in \mathbb{N}.\ \alpha \cdot \widehat{j} = \widehat{0}\}\ , \\ \top & = & \{\alpha \mid \exists t.\ (\forall j < t.\ \alpha \cdot \widehat{j} = \widehat{0}) \wedge \alpha \cdot \widehat{t} = \widehat{1}\}\ . \end{array}$$

Definition 4 (Continuity) *By the* Continuity Axiom *for A we shall mean the statement: For any $F \in A(\overline{2})$, if $F \cdot \alpha = \widehat{n}$ for all $\alpha \in \bot$, then $F \cdot \alpha = \widehat{n}$ for some $\alpha \in \top$.*

The Continuity Axiom is clearly satisfied by a wide range of models that have some notion of continuity built into their definition: for instance, by CPO models such as the Scott, stable and strongly stable models. It is also easily seen to be satisfied by models based on some evidently finitary concept of computation, such as typical game models for simply-typed languages. Finally, it is satisfied in typical 'effective' models thanks to the undecidability of the halting problem; this will be explained more fully in Section 2.7.

Of course, there are models that violate the Continuity Axiom because they contain genuinely discontinuous second-order functionals such as $\exists_{\mathbb{N}}$; the results we have in mind cannot be expected to hold for these models. (The dichotomy between 'continuous' and 'discontinuous' models is explored in detail in [16, Section 5.3].) However, there are other, intuitively 'continuous' models that do not satisfy the axiom in the above form. Perhaps most significantly, there are order-theoretic models in which distinct numerals $\widehat{n}, \widehat{m}$ may have an upper bound $\widehat{n} \sqcup \widehat{m}$: we can think of such an element as representing some kind of non-deterministic computation that might yield either m or n as its result. Examples of such models include Scott's $\mathcal{P}\omega$ and other lattice-theoretic models, as well as syntactic models for various non-deterministic languages; these models do not satisfy our axiom, because it is easy to construct an operation F such that $F(\alpha) = \widehat{n}$ whenever $\alpha \in \bot$, but $F(\alpha) = \widehat{n} \sqcup \widehat{m}$ whenever $\alpha \in \top$. In Section 6 we shall present a slightly modified version of our framework including an alternative Continuity Axiom that allows us to develop much of our theory in typical non-deterministic settings.

A further caveat is that besides the idea of continuity, our axiom also builds in the idea that each $\widehat{n}$ is in some way a 'compact' element of $A(\mathtt{N})$. To illustrate this, consider the model $\Lambda^0/\mathcal{H}^*$ of closed untyped λ-terms modulo Nakajima tree equality. This is a CPO model, and each element $\widehat{n}$ is non-compact in the sense that it is a supremum of a chain of elements strictly below it. It is thus possible to construct an element F such that $F \cdot \alpha = \widehat{n}$ for all $\alpha \in \bot$, but $F \cdot \alpha$ is strictly below $\widehat{n}$ for all $\alpha \in \top$. A somewhat similar problem also arises for Kleene's second model K_2, however, in this case, the refinements we shall introduce in Section 6 will allow this model to be accommodated naturally within our theory.

The following notation will be useful in connection with the Continuity Axiom. Let $t : \top \cup \bot \to \mathbb{N}_\infty$ be defined by

$$t(\alpha) \;=\; \mathit{min}\; j.\; f(\alpha) > 0$$

so that $t(\alpha) < \infty$ for all $\alpha \in \top$, and $t(\alpha) = \infty$ for all $\alpha \in \bot$. Intuitively, if we think of α as a process that emits zeros for as long as it is 'running', then $t(\alpha)$ gives the 'halting time' for α. For $\alpha \in \top$, the above formula gives us a way of actually computing $t(\alpha)$ within A: we write $\widehat{t} \in A(\overline{2})$ for some element such that $\alpha \in \top$ implies $\widehat{t} \cdot \alpha = \widehat{t(\alpha)}$.

We conclude this subsection with a small but important consequence of the Continuity Axiom. (A much more substantial consequence of Continuity alone will be presented as Theorem 36.)

Proposition 5 *Divergent computations exist: there exists* $\mathit{div} : \overline{1}$ *with* $\mathit{div} \cdot \widehat{0} \Uparrow$.

Proof Suppose there are no $F : \sigma \to \mathbb{N}$ and $f : \sigma$ with $F \cdot f \Uparrow$. Then $\mathit{min} \cdot f$ has a numeral value for every $f \in N^N$. Now consider $H = \lambda^* f.\ \mathit{if}\ f(\mathit{min}\ f) = \widehat{0}\ \mathit{then}\ \widehat{0}\ \mathit{else}\ \widehat{1}$. Then for all $f \in N^N$ we have $H \cdot f = \widehat{0}$ if $f \in \top$, and $H \cdot f = \widehat{1}$ if $f \in \bot$, which contradicts the Continuity Axiom.

Once we have found any F, f with $F \cdot f \Uparrow$, we may take $\mathit{div} = \lambda^* x.\, Ff = s(kF)(kf)$. □

2.3 Enumeration

Our second axiom can be seen as adapting a familiar fact from basic computability theory: namely, that if $T \subseteq \mathbb{N}$ is the range of a partial computable function and is inhabited, then T is also the range of a total computable function. To articulate this in our general setting, the following notions will be helpful:

Definition 6 *(i) Given any* $f \in A(\overline{1})$, *the* proper range *of* f *is the set*

$$\{m \in \mathbb{N} \mid \exists n \in \mathbb{N}.\, f \cdot \widehat{n} = \widehat{m}\}\ .$$

(ii) For any $g : \mathbb{N} \to \mathbb{N}$, *the* offset range *of* g *is the set*

$$\{m \in \mathbb{N} \mid \exists n \in \mathbb{N}.\, g(n) = m + 1\}\ .$$

The significance of the offset range is simply that it allows us to formulate the property of interest without treating the empty set as an exceptional case, allowing for a more constructive development of our theory. Readers with no scruples regarding classical reasoning may as well work with the ordinary range of g rather than the offset one.

A mild variant of the abovementioned range property says that for any partial computable $f : \mathbb{N} \rightharpoonup \mathbb{N}$ we can computably enumerate all pairs $\langle n, f(n) \rangle$ where $f(n)\downarrow$. We adapt this to the setting of a general model A as follows.

Definition 7 (Enumeration) *By the* Enumeration Axiom *for A we shall mean the statement: For every $f \in A(\overline{1})$, there exists $g \in \Delta_A$ whose offset range is the set of $\langle n, m \rangle$ such that $f \cdot \widehat{n} = \widehat{m}$. In other words, for all f there exists g such that for all $n, m \in \mathbb{N}$ we have*

$$\exists i.\, g(i) = \langle n, m \rangle + 1 \quad \textit{iff} \quad f \cdot \widehat{n} = \widehat{m} \; .$$

The Enumeration Axiom holds in all natural models of which we are aware. For instance, in any model in which Δ is the full set $\mathbb{N}^{\mathbb{N}}$, the axiom holds trivially; in models of an 'effective' character it holds by virtue of the familiar interleaving argument from basic computability theory. Further motivation for the axiom may be gleaned from the easy proof (in Section 3 below) that Enumeration by itself suffices to characterize the *1-spectrum* of A completely in terms of Δ_A.

One might also consider a more 'uniform' version of Enumeration which asserts that the passage from f to g can itself be effected within A via some $\Psi \in A(\overline{1} \to \overline{1})$; we may call this the *Computable Enumeration Axiom*. However, we do not adopt this as a basic axiom since relatively few models of interest satisfy it. Those that do are typically highly intensional models that support 'parallel computations via interleaving', such as Kleene's K_1 and K_2, or models for languages with timeout: see e.g. [16, Section 12.2].[3]

It is also worth noting that the combination of Continuity, P-Normalizability and Computable Enumeration is inconsistent (the proof of this offers an interesting exercise). On the other hand, Continuity, T-Normalizability and Computable Enumeration are all satisfied by Kleene's K_2.

It may come as a surprise that something as flagrantly non-uniform as our Enumeration Axiom serves our purposes well. However, it is essential that *some* such non-uniformity be present in our axioms, since the main results we are hoping to prove (e.g. every representable operation of a certain type has a 'graph') typically do not themselves hold computably within A.[4] One useful intuition is that our axioms trying to capture aspects of what *would* be true if it admitted, for instance, a well-behaved simulation in K_1. As discussed in the Introduction, we are asking what trace of the existence of such a simulation would visible within A itself, and the Enumeration Axiom can certainly be seen as such a property. One may thus informally think of the passage from f to g as being 'secretly effective' (at the level of some putative K_1-realization), though such effectivity is not typically manifest within A itself.

2.4 Normalizability

The next pair of axioms represent pleasing properties that hold in some but not all models of interest. We shall therefore be exploring both how far we can get without them, and how much further we can get with them. Informally, these axioms say that within A we may select a *canonical* choice of realizer for each representable function from $\mathbb{N}$ to $\mathbb{N}$. We shall present this principle in two flavours: a weaker one for total functions $\mathbb{N} \to \mathbb{N}$ and a stronger one for partial functions $\mathbb{N} \rightharpoonup \mathbb{N}$.

[3]The Computable Enumeration Axiom is also trivially satisfied by models with a 'full set-theoretic' flavour such as Platek's monotone model, but these models lie outside the intended scope of our investigation and are ruled out by the Continuity Axiom below.

[4]Our Continuity Axiom also lacks this uniformity; however, it is an interesting exercise is to show that Continuity plus Computable Enumeration would imply an A-computable version of Continuity.

Definition 8 (Normalizability) *(i) A* T-normalizer *for A is an element $norm : \overline{1} \to \overline{1}$ such that*

- *for any $g \in N^N$ and $n \in N$ we have $norm \cdot g \cdot n = g \cdot n$,*
- *if $g, g' \in N^N$ represent the same function $\mathbb{N} \to \mathbb{N}$, then $norm \cdot g = norm \cdot g'$.*

(ii) A P-normalizer *for A is an element $norm : \overline{1} \to \overline{1}$ such that*

- *for any $g \in A(\overline{1})$ and $n, m \in N$, we have $norm \cdot g \downarrow$, and $norm \cdot g \cdot n = m$ iff $g \cdot n = m$,*
- *if $g, g' \in A(\overline{1})$ and $g(n) = m$ iff $g'(n) = m$ for all $n, m \in N$, then $norm \cdot g = norm \cdot g'$.*

(iii) By the T-Normalizability Axiom *for A we shall mean the statement: A has a T-normalizer. The* P-Normalizability Axiom *will be the statement: A has a P-normalizer.*

Clearly any P-normalizer is a T-normalizer. T-Normalizability was one of the conditions on continuous models used in [14], while the P-Normalizability axiom was employed in [16, Section 7.1].

In many models, a P-normalizer is trivially available. As a typical example, in a domain-theoretic model in which $A(\mathtt{N}) = \mathbb{N}_\perp$ and $A(\overline{1})$ is the set of monotone functions $\mathbb{N}_\perp \to \mathbb{N}_\perp$, there will be two elements of $A(\overline{1})$ representing the constant zero function $\mathbb{N} \to \mathbb{N}$, but a suitable element *norm* may be defined by

$$norm(g)(n) \;=\; g(n) \;, \qquad norm(g)(\perp) \;=\; \perp .$$

Likewise, in a typical game model, there will be two strategies for the constant zero function — one that requests the value of its argument and one that does not — but a normalizing strategy can be used to transform any strategy of type $\mathbb{N} \to \mathbb{N}$ into the one that requests its argument.

Most models with a T-normalizer will naturally possess a P-normalizer as well, though there are a few exceptions. Notably, Kleene's second model K_2 and the λ-term model $\Lambda^0/\mathcal{B}$ modulo Böhm tree equality both satisfy T-Normalizability but not P-Normalizability.

In general, a model will have a T-normalizer provided it is 'not too fine-grained': the kinds of models that do not are highly intensional models such as Kleene's first model K_1 or the term model for PCF+timeout.

2.5 Collection

Our discussion of Enumeration introduced the idea of looking for consequences within A itself of the existence of a well-behaved simulation in K_1 (or similar), and our next axiom represents a more radical application of this idea. Although its formulation is inspired by 'effective' models, it is in fact valid for all known models of interest.

To motivate the axiom, suppose that A were realizable over K_1, and that moreover for some type σ, the set $R_\sigma \subseteq \mathbb{N}$ of all K_1-realizers for elements of $A(\sigma)$ were c.e. (This is indeed the case for practically all known effective models with the exception of K_2^{eff}.) It may or may not be the case that the function $\mathbb{N} \to A(\sigma)$ induced by an enumeration of R_σ is representable within A itself, but in any case, the image of any map $\Phi \in A(\sigma \to \mathtt{N})$ will again be c.e. at the level of realizers. Assuming that the set of K_1-realizers for $\widehat{n}$ is semidecidable uniformly in n, it follows that the proper range of Φ is c.e. Abstracting out this property, and generalizing to multi-argument functions, we obtain:

Definition 9 (Collection) *By the* Collection Axiom *for A, we shall mean the statement: For any types $\sigma_0, \ldots, \sigma_{r-1}$ and any $\Phi \in A(\sigma_0 \to \cdots \to \sigma_{r-1} \to \mathtt{N})$, there exists $f \in A(\overline{1})$ such that $f \upharpoonright_N$ has the same proper range as Φ (i.e.* $\mathrm{ran}(f \upharpoonright_N) \cap N = \mathrm{ran}(\Phi) \cap N$*).*

To be more specific, the assertion that the above holds for some particular list of types $\sigma_0, \ldots, \sigma_{r-1}$ may be referred to as the $\sigma_0, \ldots, \sigma_{r-1}$-*Collection Axiom.*

As already explained, Collection holds in most naturally occurring effective models (we will spell out the above argument more formally in Section 2.7). However, it also holds even in uncountable models A provided they contain a 'dense' effective submodel. Finally, it holds even in the anomalous case of K_2 and K_2^{eff} since it is easy to find a dense c.e. subset of these models.

Together with Enumeration, Collection serves as a way of gathering together some computably generated set of numerals into a computable listing within A. This will turn out to be useful in settings where Normalization is not available and Enumeration alone is not sufficient for the purpose. Although Collection and Enumeration will in practice be used in conjunction, we keep the two axioms separate as they appear to represent distinct conceptual ingredients.

We also mention the following stronger and more subtle axiom, which is not used within the present paper but which we expect to play a role at higher types:

Definition 10 (Dependent Collection) *By the* Dependent Collection Axiom *for A, we shall mean the statement: For any $\Psi \in A(\mathtt{N} \to \sigma_0 \to \cdots \to \sigma_{r-1} \to \mathtt{N})$, there exists $f \in A(\overline{1})$ such that*

1. *for any $\vec{x} \in A(\vec{\sigma})$ there exists n such that $f \cdot n \cong \Psi \cdot \widehat{n} \cdot \vec{x}$, and*
2. *for any $n \in \mathbb{N}$ there exists $\vec{x} \in A(\vec{\sigma})$ with $f \cdot n \cong \Psi \cdot \widehat{n} \cdot \vec{x}$.*

Again, this clearly holds in well-behaved effective models: focussing on the unary case, if $R_\sigma \subseteq N$ is as before, $h \in A(\overline{1})$ represents a computable enumeration of R_σ and $p : \mathbb{N} \to \mathbb{N}$ is a computable function tracking Ψ, then $\lambda^* n.\, p(n, h(n))$ provides a suitable element f. We therefore view Dependent Collection as a variant of Collection with a diagonal flavour. However, it is unclear at present whether Dependent Collection holds in any (or most) *uncountable* models of interest.

2.6 Restriction

Many functional programming languages include a 'sequencing' construct $-;-$, where the meaning of $M; N$ is 'evaluate M, discard the result, then evaluate N and return the result'. The effect of this is that the evaluation of $M; N$ will succeed only if the evaluations of both M and N succeed. The Restriction Axiom is a simple and natural one, saying simply that such an operator is available in our model.

Definition 11 (Restriction) *By the* Restriction Axiom *for A we will mean the statement: A contains an (infix)* restriction operator $\upharpoonright : \mathtt{N}^2 \to \mathtt{N}$ *with the properties that*

$$a \upharpoonright b \Downarrow \quad \textit{iff } a, b \in N \,, \qquad \widehat{n} \upharpoonright \widehat{m} \;=\; \widehat{n} \quad \textit{for all } n, m \in \mathbb{N} \,.$$

In the great majority of natural models, the existence of a restriction operator is completely trivial. However, in a few cases the axiom is surprisingly problematic: for instance, in the untyped λ-term models Λ^0/β or $\Lambda^0/\mathcal{B}$, it appears that restriction is not available. The Restriction Axiom will be necessary for the development of our theory in the non-normalizable case (Section 4.2), but we will see that it can be dispensed with in the presence of T-Normalizability (Section 4.3), allowing models such as $\Lambda^0/\mathcal{B}$ to be accounted for.

As a typical illustration of the use of restriction, suppose we require an element *eqtest* : $\mathtt{N}^2 \to \mathtt{N}$ with the property

$$eqtest\ \widehat{n}\ a \Downarrow \quad \text{iff} \quad a = n\ .$$

We can define such an element as follows:

$$eqtest \;=\; \lambda^* n.\lambda^* a.\ (if\ (eq\, n\, a)\ then\ \widehat{0}\ else\ div\, \widehat{0}) \;\upharpoonright\; a\ ,$$

where *eq* implements equality testing on numerals and *div* is as in Proposition 5. Note that without the restriction to a, we could not guarantee that that the if-expression never spuriously returns a numeral value on some non-numerals a.

Using a similar idea, we can now improve on the construction of Proposition 5 as follows:

Lemma 12 *Assume A satisfies Continuity and Restriction. Then there is a element $min' \in A(\overline{2})$ that satisfies the usual condition for minimization and an additional one:*

1. *If $f \cdot \widehat{n} = \widehat{0}$ for all $n < r$ and $f \cdot \widehat{r} = \widehat{m+1}$ for some m, then $min' \cdot f = \widehat{r}$.*
2. *If $f \cdot \widehat{n} = \widehat{0}$ for all n, then $min' \cdot f \Uparrow$.*

Proof Define $min' = \lambda^* f.\ (if\ f(min f) > 0\ then\ min\ f\ else\ div\ \widehat{0}) \upharpoonright (min\ f)$. $\Box$

Corollary 13 *Assume A satisfies Continuity and Restriction. Then any partial computable $f : \mathbb{N} \rightharpoonup \mathbb{N}$ is* strongly representable *in A: there exists $\widehat{f} \in A(\overline{1})$ such that $f(n) = m$ implies $\widehat{f} \cdot \widehat{n} = \widehat{m}$, and $f(n) \uparrow$ implies $\widehat{f} \cdot \widehat{n} \Uparrow$.*

Proof Using the Kleene normal form theorem, take e such that $f(n) = U(min\ y.\, T(e, n, y))$ for each n. Take $\widehat{T}, \widehat{U} \in A$ representing the primitive recursive functions T, U, and define

$$\begin{aligned} h &= \lambda^* n.\ min'(\lambda^* y.\, \widehat{T}(\widehat{e}, n, y))\ , \\ \widehat{f} &= \lambda^* n.\ (\widehat{U}\,(hn)) \upharpoonright (hn) \end{aligned}$$

with min' as in Lemma 12. $\Box$

Lemma 12 also yields the following key property:

Theorem 14 (Non-Invertibility) *Assume A satisfies Continuity and Restriction.*

(i) There is no element $\neg \in A(\overline{1} \to \overline{1})$ such that

$$\alpha \in \top \;\Rightarrow\; \neg \cdot \alpha \in \bot\ , \qquad \alpha \in \bot \;\Rightarrow\; \neg \cdot \alpha \in \top\ .$$

(ii) Let $\top_0 = \{f \in A(\overline{1}) \mid f \cdot \widehat{0} \Downarrow\}$ and $\bot_0 = \{f \in A(\overline{1}) \mid f \cdot \widehat{0} \Uparrow\}$. Then there is no element $\neg_0$ mapping $\top_0$ into $\bot_0$ and vice versa.

PROOF (i) The element min' from the previous lemma satisfies $min' \cdot \alpha \Downarrow$ if $\alpha \in \top$ and $min' \cdot \alpha \Uparrow$ if $\alpha \in \bot$. So if $\neg$ existed as above, then $\lambda^\alpha.\ \widehat{0} \upharpoonright (min'(\neg\alpha))$ would map all $\alpha \in \bot$ to $\widehat{0}$ and would diverge for all $\alpha \in \top$, contradicting the Continuity Axiom.

(ii) Note that $\lambda^* f.\lambda^* x.min'\, f$ maps $\top$ into $\top_0$ and $\bot$ into $\bot_0$. So if $\neg_0$ existed, then $\lambda^* f.\ \widehat{0} \upharpoonright (\neg_0\, (\lambda^* x.min'\, f)\, \widehat{0})$ would map $\bot$ to $\widehat{0}$ and $\top$ to divergence, again contradicting Continuity. □

The following consequence of part (ii) above is sometimes useful. If $x, y \in A(\sigma)$, let us write $x \ll y$ if there exists $F \in A(\overline{1} \to \sigma)$ mapping all of $\bot_0$ to x and all of $\top_0$ to y. We refer to $\ll$ as the *link relation*; note that it is not in general a partial order. As we shall see in Section 4.4, instances of the link relation arise quite naturally in the presence of the P-Normalizability Axiom.

Lemma 15 *Suppose $x \ll y$. If $g \cdot x = \widehat{n}$, then $g \cdot y = \widehat{n}$.*

PROOF Suppose F maps $\bot_0$ to x and $\top_0$ to y, and take *eqtest* as above. If $g \cdot x = \widehat{n}$ but $g \cdot y \neq \widehat{n}$, then clearly $\lambda^* f.\lambda^* z.\ eqtest\ \widehat{n}\ (g(F(f)))$ interchanges $\bot_0$ and $\top_0$, contrary to Theorem 14(ii). □

Finally, for certain specialized purposes in Section 4.2, we shall need a more powerful operator which allows the conjunction of a dynamically chosen list of restriction conditions:

Definition 16 (Iterated Restriction) *By the* Restriction Axiom *for A we will mean the statement: A contains an operator $\upharpoonright\upharpoonright : \mathtt{N} \to \mathtt{N} \to (\mathtt{N} \to \mathtt{N}) \to \mathtt{N}$ such that:*

- *$a \upharpoonright\upharpoonright_{\widehat{t}}\, f \Downarrow$ iff $a \Downarrow$ and $f \cdot i \Downarrow$ for all $i < t$,*
- *in this case, $a \upharpoonright\upharpoonright_{\widehat{t}}\, f = a$.*

Clearly Iterated Restriction implies Restriction, and all models we know that satisfy Restriction also satisfy Iterated Restriction. However, to construct an iterated restrictor from an ordinary one requires a slightly more powerful iteration principle than the operation *iter* of Definition 3, since we potentially need to carry round a non-numeral as data in addition to the loop counter. Since we have no other uses for this stronger kind of iteration, it seems simplest to posit Iterated Restriction directly as a possible axiom.

As a sample application, in the presence of Iterated Restriction we can improve our minimization operator yet further: there is now an element min'' such that for any f and $r \in N$, we have $min'' \cdot f = \widehat{r}$ if and only if $f \cdot \widehat{n} = \widehat{0}$ for all $n < r$ and $f \cdot \widehat{r} = \widehat{m+1}$ for some m (so that for instance $min'' \cdot f \Uparrow$ if $f \cdot \widehat{0} \Uparrow$).

2.7 Relationship to extrinsic conditions involving simulations

Before we proceed to develop our mathematical theory on the basis of the above axioms, we pause to consider how our 'intrinsic' axiomatic approach relates to the more 'extrinsic' approach involving simulations. In particular, we shall clarify which of our axioms are readily implied by the existence of a well-behaved simulation in K_2 or K_1. This makes the connection with our previous work in [14], and indeed enables us to compare the levels of generality of the two approaches. That said, the material in this section is not really required for the rest of the paper, and since it largely concerns somewhat artificial conditions from [14] which our present approach has the merit of avoiding, the reader may prefer to skip it.

We assume familiarity here with Kleene's untyped PCAs K_1 and K_2 (here regarded as TPCAs), and with the general concept of an *applicative simulation* between TPCAs (see [16, Section 3.3]). The idea in [14] was that a TPCA A with numerals could be regarded as *continuous* if it admitted a certain kind of simulation $\gamma : A \longrightarrow\!\!\triangleright K_2$, and *effective* if it admitted a simulation $\gamma : A \longrightarrow\!\!\triangleright K_1$. We shall consider the effective case first as it is somewhat easier to explain.

First, we required that $\gamma : A \longrightarrow\!\!\triangleright K_1$ 'respects numerals': if $\widehat{0}, \widehat{1}, \ldots$ is a system of numerals in A, there is an element $d \in K_1$ such that if $a \Vdash^\gamma \widehat{n}$ then $d \cdot a = n$. (The existence of a realizer that translates in the other direction is immediate.) An applicative simulation γ with this property was called a *realization* in [14]. Clearly, this condition implies that the functions $\mathbb{N} \to \mathbb{N}$ representable within $A(\overline{1})$ are all Turing computable, a property one would certainly expect of an 'effective' TPCA. Conversely, in both [14] and the present paper, the computational hypotheses on A are strong enough to ensure that all computable functions $\mathbb{N} \to \mathbb{N}$ are A-representable.

As hypotheses for our ubiquity theorems in [14], we also formulated the following conditions on effective models:

(A) For any type σ, there is an element $\alpha_\sigma \in K_1$ that tracks application of elements of $A(\sigma \to \mathbb{N})$ in the following sense: if $m \Vdash^\gamma_{\sigma \to \mathtt{N}} f$ and $n \Vdash^\gamma_\sigma x$, then $\alpha_\sigma \cdot m \cdot n \Vdash^\gamma_{\mathbb{N}} y \in A(\mathtt{N})$ iff $f \cdot x = y$.

(B') The relation $m \cdot^\gamma_N \widehat{n}$ is c.e. in m, n.

Condition (A) is a very mild requirement that holds in all effective models have ever considered. Condition (B') is also a natural requirement, in that we expect to be able to recognize when a computation yields the numerical result n. However, this condition does not hold in Scott's $\mathcal{P}\omega^{\mathrm{eff}}$ and similar models of a non-deterministic flavour (cf. the discussion in Section 2.2), and in [14] we did not have a way of dealing with such models.

One of the major results of [14] was that in any effective model A satisfying conditions (A) and (B)', the type structure of hereditarily extensional total functionals over $\mathbb{N}$ (obtained as the *extensional collapse* with respect to the set of numerals) coincides with the classical type structure HEO of *Hereditarily Effective Operations*. Actually, this result was obtained under a marginally weaker version of (B)':

(B) The relation $m \cdot^\gamma_N \widehat{0}$ is c.e.

However, it would seem to be a kind of accident that this was sufficient, and (B') (which featured elsewhere in [14]) would appear to be the more natural condition.

We may relate these conditions to our present axioms as follows:

Proposition 17 *Any effective model satisfying (A) and (B') above satisfies Continuity.*

PROOF For Continuity, suppose $\top, \bot \subseteq A(\overline{1})$ are as defined in Section 2.2. Let $H \subseteq \mathbb{N}$ denote the halting set. We shall construct a computable function $\theta : \mathbb{N} \to \mathbb{N}$ such that

$n \in H$ implies $\theta(n) \Vdash^\gamma \alpha$ for some $\alpha \in \top$, and $n \notin H$ implies $\theta(n) \Vdash^\gamma \alpha$ for some $\alpha \in \bot$. First, given n it is straightforward to compute a K_1-realizer $\delta(n)$ such that

$$\begin{array}{lcl} n \in H & \Rightarrow & \forall j.\ \delta(n) \cdot j = 0\ , \\ n \notin H & \Rightarrow & \exists j.\ (\forall j < t.\ \delta(n) \cdot j = 0) \wedge \delta(n) \cdot t = 1\ . \end{array}$$

Next, using the realizer d from the definition of effective TPCA, we may transform $\delta(n)$ into a realizer for the numeral $\widehat{\delta(n)} \in A(\mathtt{N})$. Finally, if $app \in A(\mathtt{N} \to \mathtt{N} \to \mathtt{N})$ represents Kleene application, then clearly $n \in H$ implies $app \cdot \widehat{\delta(n)} \in \top$, and $n \notin H$ implies $app \cdot \widehat{\delta(n)} \in \bot$; this computation may also be simulated at the K_1 level via γ, so that from a realizer for $\widehat{\delta(n)}$ we may compute a realizer for an element of $\top$ or $\bot$ as appropriate. Combining these constructions gives the required θ.

To verify the Continuity Axiom, suppose $F \in A(\overline{2})$ satisfies $F \cdot \alpha = \widehat{p}$ for all $\alpha \in \bot$. Using γ, take $\phi : \mathbb{N} \rightharpoonup \mathbb{N}$ computable such that $\phi(m) \Vdash^\gamma F(\alpha)$ whenever $m \Vdash^\gamma \alpha$ and $F(\alpha) \downarrow$; and using condition (B'), take $\epsilon_p : \mathbb{N} \rightharpoonup \mathbb{N}$ computable such that $\epsilon_p(m) \downarrow$ iff $m \Vdash^\gamma \widehat{p}$. We then have that $\epsilon_p(\phi(\delta(n))) \downarrow$ for all $n \notin H$. By the undecidability of the halting problem, there is therefore some $n \in H$ such that $\epsilon_p(\phi(\delta(n))) \downarrow$, i.e. $\phi(\delta(n)) \Vdash^\gamma \widehat{p}$. But now $\delta(n)$ realizes some $\alpha \in \top$, so condition (A) implies that $F(\alpha) = \widehat{p}$ as required. $\square$

Proposition 18 *Any effective model satisfying (A) and (B') satisfies Enumeration.*

Proof Suppose $f \in A(\overline{1})$, and use γ along with condition (A) to obtain a computable $\psi : \mathbb{N} \rightharpoonup \mathbb{N}$ such that if $m \Vdash^\gamma \widehat{n}$ then $\psi(m) \Vdash^\gamma \widehat{p}$ iff $f \cdot \widehat{n} = \widehat{p}$. Using the realizer d to encode natural numbers as realizers for numerals, and condition (B') to decode such realizers into natural numbers again, we obtain a computable $\chi : \mathbb{N} \rightharpoonup \mathbb{N}$ such that $\chi(n) = p$ iff $f(\widehat{n}) = \widehat{p}$. By basic computability theory, there is a total computable function g whose offset range is precisely the set of such pairs $\langle n, p \rangle$. Since A represents all computable functions $\mathbb{N} \to \mathbb{N}$, this g is present in Δ_A. (We may, if desired, effectively compute a γ-realizer for an A-representative of g from a γ-realizer for F, although this is not required here.) $\square$

Our conditions on effective models imply nothing as regards Normalizability — indeed it is explicit in [14] that the results for effective models apply even to highly intensional models. They also do not imply Restriction, although we view this as a relatively mild condition which we are happy to adopt when necessary.

It remains to consider Collection. We here firm up the argument already sketched in Section 2.5 to the effect that Collection holds in almost all natural effective models. Let us introduce the following condition on models:

(D) For each type σ, the set of n such that $n \Vdash^\gamma_\sigma x$ for some $x \in A(\sigma)$ is c.e.

This condition did not feature in [14], but it appears to hold for all the effective models we have in mind, with the somewhat anomalous exception of K_2^{eff} (for which Collection can be seen to hold anyway). In particular, if A arises as a term model for some programming language, and γ is given by Gödel numbering of terms, then condition (D) simply amounts to saying that the terms of the language are computably enumerable. Moreover, we have:

Proposition 19 *Any effective model satisfying (A), (B') and (D) and the Restriction Axiom also satisfies σ-Collection for any type σ.*

PROOF Given $\Phi \in A(\sigma \to \mathbb{N})$, by condition (A) we may again take a computable $\psi : \mathbb{N} \rightharpoonup \mathbb{N}$ such that if $m \Vdash^\gamma x \in A(\sigma)$ then $\psi(m) \Vdash^\gamma \widehat{n}$ iff $\Phi \cdot x = \widehat{n}$. Using condition (D) to enumerate all such realizers m, and (B') to decode the resulting realizers for $\widehat{n}$, we obtain a partial computable function θ whose range is the set of all n such that $\Phi \cdot x = \widehat{n}$ for some $x \in A(\sigma)$.

Finally, we note that our model satisfies Continuity by Proposition 17, and Restriction by hypothesis; hence by Corollary 13, θ is strongly representable in A, and this gives us $f \in A(\overline{1})$ with $\mathrm{ran}(f \upharpoonright_N) \cap N = \mathrm{ran}(\Phi) \cap N$. $\Box$

An obvious generalization of condition (A) to types $\sigma_0 \to \cdots \to \sigma_{r-1} \to \mathtt{N}$ likewise yields $\sigma_0, \ldots, \sigma_{r-1}$-Collection for all $\sigma_0, \ldots, \sigma_{r-1}$.

We now turn our attention to continuous models. Here we consider a model A endowed with an applicative simulation $\gamma : A \longrightarrow\!\!\triangleright\, K_2$. As in the effective case, we shall require that γ 'respects numerals', but this condition must now be formulated in terms of some standard choice of numerals $\widetilde{0}, \widetilde{1}, \ldots$ for K_2 as well as a choice $\widehat{0}, \widehat{1}, \ldots$ for A. That is, we require that there is an element $d \in K_2$ such that if $a \Vdash^\gamma \widehat{n}$ then $d \cdot a = \widetilde{n}$. In this situation we say that A is a *continuous* TPCA with realization γ.

We say that A is *full continuous* with respect to γ if every set-theoretic function $f : \mathbb{N} \to \mathbb{N}$ is represented in $A(\overline{1})$, and moreover there is an element $h \in K_2$ such that for any $f : \mathbb{N} \to \mathbb{N}$, $h \cdot f \Vdash^\gamma \widehat{f}$ for some $\widehat{f} \in A(\overline{1})$ representing f. The situation treated in [14] was that of a full continuous model satisfying the following special conditions:

(AB) For any type σ, any $f \in A(\sigma \to \mathtt{N})$ and any $p \in \mathbb{N}$, there is an open subset $U_p \subseteq \mathbb{N}^{\mathbb{N}}$ such that if $b \Vdash^\gamma_\sigma x$ then $f \cdot x = \widehat{p}$ iff $b \in U_p$.

(C) A satisfies T-Normalizability.

Condition (AB) may appear rather obscure, but its purpose is to capture the import of the above conditions (A) and (B) in the K_2 setting via a single condition. This slightly strange approach is necessitated by the annoying fact that the definedness of application in K_2 is not semidecidable. Condition (C) was a hypothesis we were unable to dispense with in the continuous setting, but managed to avoid in the effective setting. Under these conditions, it was shown in [14] that the hereditarily extensional total functionals arising from A are precisely the Kleene-Kreisel continuous functionals.

Proposition 20 *Any full continuous model satisfying (AB) satisfies Continuity.*

PROOF Let $a_\infty \in K_2$ be the constant zero function, and a_n be the function mapping i to 0 if $i \leq n$ and 1 otherwise. The realizer h from the definition of full continuity then maps a_∞ to a realizer for an element of $\bot$, and each a_n to a realizer for an element of $\top$. To verify Continuity, suppose $F \in A(\overline{2})$ maps every element of $\bot$ to $\widehat{p}$, and take $\widehat{F} \in K_2$ such that whenever $\widehat{f} \Vdash^\gamma_{\overline{1}} \widehat{f}$ and $F \cdot f\downarrow$, we have $\widehat{F} \cdot \widehat{f} \Vdash^\gamma_{\mathtt{N}} F \cdot f$. Then $\widehat{F} \cdot (h \cdot a_\infty) \Vdash \widehat{p}$. Let U_p be the open set given by condition (AB) so that if $b \Vdash^\gamma_{\overline{1}} \alpha$ then $F \cdot \alpha = \widehat{p}$ iff $b \in U_p$. Then $h \cdot a_\infty \in U_p$. Moreover, $h \cdot a_\infty$ is the limit of the $h \cdot a_n$, so $h \cdot a_n \in U_p$ for some n. But now $h \cdot a_n$ realizes some $\alpha \in \top$, where $F \cdot \alpha = \widehat{p}$. $\Box$

In the case of full continuous models, both Enumeration and Collection are immediate from the fact that *every* $f : \mathbb{N} \to \mathbb{N}$ is represented in A. For Enumeration, the construction of the required $g \in \Delta$ from $f \in A(\overline{1})$ can also be carried out 'effectively in K_2' if desired;

it appears that a similarly ‘constructive’ treatment is also possible using a suitable K_2 analogue of condition (D), though we shall not pursue this here.

As before, our conditions do not imply Restriction, though as commented earlier, we find that we have little need for it in the presence of T-Normalizability, which is here explicitly built in by condition (C).

Finally, in the light of these observations, we may comment on how our coverage of models via our present approach compares with that in [14] (with the caveat that we have not yet carried through our axiomatic proofs of ubiquity beyond level 2).

The first point to make is that whereas in [14] the ‘full continuous’ and ‘effective’ situations were treated separately, there is no such dichotomy in our present approach. All our main proofs will work uniformly regardless of what the A-representable functions $\mathbb{N} \to \mathbb{N}$ are (although the set of such functions does have a key role to play). They therefore apply to many full continuous models as well as many effective ones, and indeed to other kinds of models such as those based on hyperarithmeticity. Nevertheless, there is still something to say about the various axiom combinations we shall consider and the extent to which they cover the full continuous and effective settings as defined above.

In Section 4.3, we shall analyse the situation in the presence of Continuity, Enumeration and T-Normalizability. As shown by the above discussion, this covers the whole class of full continuous models treated in [14]. Although we have so far only completed the analysis at type level 2, we fully expect that the same combination of axioms will allow the proof in [14] for the full continuous case to be adapted to our present setting.

In Section 4.2, we shall consider the situation without normalizability assumptions, but with Continuity, Enumeration, Iterated Restriction and Collection. This in practice covers the vast majority of effective models addressed by the treatment in [14]: we have seen that all of the latter satisfy Continuity and Enumeration, and also Collection under the additional condition (D). We are thus left only with the handful of models that fail to satisfy (D) and Iterated Restriction; moreover, some of the most prominent examples of such models (e.g. $\Lambda^0/\mathcal{B}$) actually satisfy T-Normalizability, so are conveniently covered by the treatment of Section 4.3. There are also some minor gains: whereas our theorems in [14] covered the (effective) term model for a version of PCF+timeout, our treatment in Section 4.2 covers both this and the ‘oracle’ version PCF^{Ω}+timeout, which we were unable to handle in [14].

As regards the extension to higher types, however, the picture is less clear in the effective setting. We are hopeful that in the absence of T-Normalizability, we will be able to carry over the proof in [14] for the effective case to the axiomatic setting of Section 4.2 (perhaps augmented by Dependent Collection or similar), although this has yet to be confirmed.

For further remarks on the prospects for extensions to higher types, see Section 8.

3 The 1-spectrum of a model

We are now ready to start proving some consequences of our axioms. We will show that in various combinations, these axioms enable us to characterize significant parts of the *spectrum* of a model A purely in terms of the set Δ_A. As a gentle first step, we shall show in this section how the Enumeration Axiom alone allows us to characterize completely the *1-spectrum* of A. In Sections 4 and 5 we will see how similar results may be obtained (much less trivially) for portions of the 2-spectrum.

It is convenient to frame our definition of the spectrum using the terminology of *modest sets*. We recall the following standard notions:

Definition 21 *Let A be any TPCA.*

(i) A modest set *X over A consists of an ordinary set $|X|$, a choice of type σ_X, and a* realizability relation *$\Vdash_X \subseteq A(\sigma_X) \times |X|$ such that:*

- *for any $x \in |X|$ there exists at least one $a \in A(\sigma_X)$ such that $a \Vdash_X x$,*
- *if $a \Vdash_X x$ and $a \Vdash_X x'$ then $x = x'$.*

(ii) If X, Y are modest sets, a morphism *$X \to Y$ is an ordinary function $|X| \to |Y|$ that is* tracked *by some $t \in A(\sigma_X \to \sigma_Y)$, in the sense that for any $x \in |X|$ and $a \in A(\sigma_X)$ we have*

$$a \Vdash_X x \implies t \cdot a \downarrow \ \wedge \ t \cdot a \Vdash_Y f(x) .$$

(iii) If X, Y are modest sets, the exponential *modest set $X \Rightarrow Y$ is defined as follows: $|X \Rightarrow Y|$ is the set of morphisms $f : X \to Y$; $\sigma_{X \Rightarrow Y} = \sigma_X \to \sigma_Y$; and $t \Vdash_{X \Rightarrow Y} f$ iff t tracks f.*

(iv) If X is a modest set and $S \subseteq X$, the (regular) subobject *of X determined by S, written $\mathsf{Sub}(X, S)$, is the modest set Y defined by: $|Y| = S$, $\sigma_Y = \sigma_X$, and $a \Vdash_Y x$ iff $x \in S$ and $a \Vdash_X x$.*

It is well-known that modest sets over A and the morphisms between them form a category $\mathcal{M}od(A)$ with a great deal of structure: for instance, it is cartesian closed (with exponentials defined as above) and has equalizers (with regular monos corresponding to the subobjects indicated above).

It is now a simple matter to define the 1-spectrum. We start with the (trivial) 0-spectrum in order to give an idea of the general construction.

Definition 22 *Let A be any TPCA with weak numerals, and let $N \subseteq A(\mathtt{N})$ denote its set of numerals.*

(i) The 0-spectrum *of A consists of simply the set $\mathbb{N}$, and the* realized 0-spectrum *of A consists of the modest set $N = (\mathbb{N}, \mathtt{N}, \Vdash_N)$ where $a \Vdash_N n$ iff $a = \widehat{n}$. (We tolerate some overloading of the symbol N here.)*

(ii) The 1-spectrum *of A is simply the mapping $\mathsf{S}^1 = \mathsf{S}^1_A$ that associates to each subset $Q \subseteq \mathbb{N}$ the set $\mathsf{S}^1(Q)$ of all morphisms from $\mathsf{Sub}(N, Q)$ to N.*

(iii) The realized 1-spectrum *of A is the mapping $\mathsf{R}^1 = \mathsf{R}^1_A$ that associates to each $Q \subseteq \mathbb{N}$ the modest set $\mathsf{R}^1(Q) = \mathsf{Sub}(N, Q) \Rightarrow N$.*

The 1-spectrum is thus a purely set-theoretic structure built up from $\mathbb{N}$ and not containing any elements from A itself; it is therefore reasonable to ask whether two models A, B share the same 1-spectrum. The realized 1-spectrum enriches the 1-spectrum with a realizability structure specific to A; note that $|\mathsf{R}^1(Q)| = \mathsf{S}^1(Q)$.

The point of the 1-spectrum is that there may be realizable functions $Q \to \mathbb{N}$ that do not extend to realizable functions $\mathbb{N} \to \mathbb{N}$. (For instance, if Q is the halting set, the function associating to each $n \in Q$ its halting time does not extend to a total computable function on $\mathbb{N}$.) So on the face of it, S_1 captures more information about computability in A than is immediately given by the set $\Delta_A \subseteq \mathbb{N}^{\mathbb{N}}$. Nonetheless, Δ_A does determine a lower bound on the contents of the 1-spectrum in the following way.

Definition 23 *Suppose $Q \subseteq \mathbb{N}$, $f : Q \to \mathbb{N}$ and $g : \mathbb{N} \to \mathbb{N}$. We say g is a* graph *of f if for every $n \in Q$ and $m \in \mathbb{N}$, $\langle n, m \rangle$ is in the range of g if and only if $m = f(n)$.*

Thus, the range of a graph of $f : Q \to \mathbb{N}$ contains all $\langle n, f(n) \rangle$ for $n \in Q$, but is allowed to contain values $\langle n, m \rangle$ where $n \notin Q$, and even values that do not code pairs at all. Note too that a graph may contain repetitions: we may have $g(i) = g(j)$ where $i \neq j$.

Proposition 24 *Suppose A is a standard model with iteration. If $f : Q \to \mathbb{N}$ has a graph within Δ_A, then $f \in \mathsf{S}_1(Q)$.*

Proof Suppose $g \in \Delta_A$ is a graph of f. Since g is itself realized by an element of N^N, the following operation $h : \mathtt{N} \to \mathtt{N}$ can be programmed within A: given an argument $n \in N$, use minimization to search for the least i such that $\mathit{fst}(g(i)) = n$; and if such an i is found, return $\mathit{snd}(g(i))$. This gives a realizer h for f, thus $f \in \mathsf{S}^1(Q)$. $\square$

The key idea is that we shall consider a model A to be 'well-behaved' if this is *all* that its 1-spectrum contains: that is, if the whole of S^1_A is determined by Δ_A in this sense.

Definition 25 *Let A be a TPCA with weak numerals. We say A is* 1-regular *if for all $Q \subseteq \mathbb{N}$, every $f \in \mathsf{S}^1_A(Q)$ has a graph in Δ_A.*

Thus, in a 1-regular model, the set Δ_A determines the whole of S^1_A by:

$$\mathsf{S}^1_A(Q) \ = \ \{f : Q \to \mathbb{N} \ \mid \ f \text{ has a graph in } \Delta_A\} \ .$$

Theorem 26 *If A satisfies Enumeration, then A is 1-regular.*

Proof Suppose A satisfies Enumeration. Suppose that $Q \subseteq N$ and that $f \in \mathsf{S}^1(Q)$ is realized by $h \in A(\mathtt{N} \to \mathtt{N})$. We wish to show that f has a graph in Δ. Applying the Enumeration Axiom to h, we obtain $g \in \Delta$ whose offset range consists exactly of all pairs $\langle n, m \rangle$ with $h \cdot \widehat{n} = \widehat{m}$. Finally, define

$$g'(n) \ = \ \mathit{if} \ n = 0 \ \mathit{then} \ \langle\rangle \ \mathit{else} \ g(n) - 1$$

and note that $g' \in \Delta$. Recalling that h realizes f, it is now easy to see that g' serves as a graph of $f : Q \to \mathbb{N}$. $\square$

Notice that the above proof would not work if our 'graphs' of f were required not to feature any pairs $\langle n, m \rangle$ where $n \notin R$. For the same reason, the converse of the above theorem does not hold: if we try to use a graph of f to construct an offset enumeration for the proper range of f, we may find that the enumeration contains spurious elements.

Although trivial in the present setting, this use of Enumeration will serve as a model for many of its uses later on.

4 The 2-spectrum below $\mathbb{N}$

We now come to the core of the paper and to our main results. We will introduce the concept of the *2-spectrum* of A, and show that under various axiom combinations, significant parts of the 2-spectrum are completely determined by Δ_A.

The following naturally extends our definition of the 1-spectrum to the next level:

Definition 27 *Let A be any TPCA with weak numerals, with $N \subseteq A(\mathtt{N})$ the set of numerals.*

(i) The 2-spectrum *of A is the mapping S^2 associating to each $Q \subseteq \mathbb{N}$ and each $R \subseteq \mathsf{S}^1(Q)$ the set $\mathsf{S}^2(Q,R)$ of all morphisms from $\mathsf{Sub}(\mathsf{R}^1(Q),R)$ to N.*

(ii) The realized 2-spectrum *of A is the mapping R^2 associating to each $Q \subseteq \mathbb{N}$ and $R \subseteq \mathsf{S}^1(Q)$ the modest set $\mathsf{R}^2(Q,R) = \mathsf{Sub}(\mathsf{R}^1(Q),R) \Rightarrow N$.*

For completeness, we note here the evident generalization to arbitrary levels, although this will not be formally required until Section 7.

Definition 28 *(i) The* realized spectrum *R_A of A is an infinite tree whose nodes are labelled with modest sets over A, constructed as follows:*

- *The root node is labelled with the modest set $\mathbb{N}$.*
- *If γ is a node labelled with a modest set X, then for each $R \subseteq |X|$ there is a branch from γ labelled by R, leading to a node labelled with the modest set $\mathsf{Sub}(X,R) \Rightarrow \mathbb{N}$.*

(ii) The spectrum *S_A of A is the tree obtained from R_A by replacing each modest set X with its underlying set $|X|$.*

(iii) The realized k-spectrum *R^k_A of A is the subtree of R_A consisting of the nodes down to level k, counting the root node as of level 0. Similarly for the k-spectrum S^k_A.*

We pause to comment briefly on what is *not* present in the 2-spectrum. Within the category $\mathcal{M}od(A)$ of modest sets on A, there are in general many more subobjects of a modest set X than those arising as $\mathsf{Sub}(X,S)$ for some $S \subseteq |X|$. For instance, we might consider a modest set Y with $|Y| \subseteq |X|$ but in which an element $y \in |Y|$ may have fewer $\Vdash_Y$ realizers than $\Vdash_X$ realizers. Subobjects of this kind are not accounted for in the spectrum, and indeed it seems that such subobjects would quickly take us into regions of $\mathcal{M}od(A)$ whose contents were highly specific to the particular choice of A (although this requires further investigation).

As in the case of the 1-spectrum, the contents of Δ_A in general provide a lower bound on the contents of S^2. For this, we need a notion of graph for second-order functionals. For later convenience, we give a definition that also makes sense for partial functionals F, although our interest at present is in total ones.

Definition 29 *(i) A number a is a* 1-code *if it is of the form $\langle\langle q_0,p_0\rangle,\ldots,\langle q_{r-1},p_{r-1}\rangle\rangle$. Such an a* matches *a function $f:\mathbb{N}\to\mathbb{N}$ if $f(q_i)=p_i$ for each $i<r$. We write U_a for the set of all functions $\mathbb{N}\to\mathbb{N}$ matched by a.*

(ii) Suppose $Q \subseteq \mathbb{N}$, $R \subseteq \mathsf{S}^1(Q)$ and $F : R \rightharpoonup \mathbb{N}$. We say $G \subseteq \mathbb{N}$ is a set graph *of F if:*

1. *For every $f \in \mathrm{dom}\, F$, G contains some $\langle a,m\rangle$ where a is a 1-code matching f.*
2. *For each $\langle a,m\rangle \in G$ where a is a 1-code, and for every $f \in R$ matched by a, we have $F(f)=m$.*
3. *For every $\langle\langle\langle q_0,p_0\rangle,\ldots,\langle q_{r-1},p_{r-1}\rangle\rangle,m\rangle \in G$, we have $q_i \in Q$ for each i.*

(iii) Under the hypotheses of (ii), we say $g:\mathbb{N}\to\mathbb{N}$ is an enumerated graph *(or just* graph*) of F if the range of g is a set graph of F.*

Once again, we allow a graph of F to contain repetitions, and to feature elements not of the specific form above, as well as elements of this form that do not 'cover' any $f \in R$. However, we do require condition 3 above to ensure that it makes sense to test a graph element against an arbitrary $f \in R$; the importance of this is made clear by the proof of the next proposition.

From here on, A will be assumed to be a standard model with iteration.

Proposition 30 *Let Q, R be as above. If $F : R \to \mathbb{N}$ has a graph within Δ_A, then $F \in \mathsf{S}^2(Q, R)$.*

PROOF Suppose g is a graph of F within Δ_A, so that g has a realizer $\widehat{g} \in N^N \subseteq A(\overline{1})$. We need to show that F exists as a morphism $\mathsf{Sub}(\mathsf{R}^1(Q), R) \to N$, that is, there is a realizer $\widehat{F} \in A(\overline{2})$ such that whenever $f \in R$ and $\widehat{f} \cdot \widehat{q} = \widehat{f(q)}$ for all $q \in Q$, we have $\widehat{F} \cdot \widehat{f} = \widehat{F(f)}$.

Informally, the algorithm for $\widehat{F}$ is as follows. Given a realizer $\widehat{f}$ for $f \in R$:

- Search the range of g looking for 'graph elements' of the form

$$\langle\langle\langle q_0, p_0\rangle, \ldots, \langle q_{r-1}, p_{r-1}\rangle\rangle, m\rangle$$

- For each such element found, test for each $i < r$ in turn whether $f(q_i) = p_i$. Each such test is guaranteed to yield a yes/no result by Condition 3 of Definition 29.
- If the graph element in question passes all these tests, stop the search and return $\widehat{m}$.

It is easy to see how this algorithm is programmable in a standard model with iteration, making use of $\widehat{g}$, and that this yields a realizer $\widehat{F}$ as required. $\square$

Remark 31 The above proof illustrates the informal style in which we shall typically present algorithms. Technically, the computation performed by $\widehat{F}$ works entirely at the level of realizers within A, but for readability we shall wherever possible couch it in terms of properties of the mathematical entities they represent, as long as the implementation in terms of realizers is clear. For example, in the above, the condition $f(q_i) = p_i$ will actually be tested by evaluating something of the form $eq \cdot (\widehat{f} \cdot \widehat{q_i}) \cdot \widehat{p_i}$. We hope that in each similar case below the intention will be sufficiently clear from the context.

Definition 32 *For Q, R as above, we say the 2-spectrum of A is* regular at Q, R *if every $F \in \mathsf{S}^2(Q, R)$ has a graph within Δ_A. We say A is* 2-regular *if the whole of its 2-spectrum is regular.*

More generally, we can consider restricted notions of regularity that apply only to certain portions of the spectrum. For instance, we will say A is *2-regular below* $\mathbb{N}$ if its 2-spectrum is regular at $\mathbb{N}, R$ for all $R \subseteq \mathsf{S}^1(\mathbb{N})$.

We will show how various combinations of axioms imply regularity for various parts of the 2-spectrum. In this section we concentrate on the 2-spectrum below $\mathbb{N}$; other subsets $Q \subseteq \mathbb{N}$ will be considered in Section 5.

As we have seen, Proposition 30 already gives us a lower bound on the contents of the 2-spectrum, so it remains to show that this is also an upper bound. Our main technology here will come from the original proof of the *Kreisel-Lacombe-Shoenfield (KLS) theorem* [13] (see also [16, Section 9.2]). In its simplest form, this theorem in effect says that in the case $A = K_1$, every modest set morphism $(N \Rightarrow N) \to N$ is continuous and indeed has a computable graph: this is tantamount to regularity at $\mathbb{N}, \Delta$. Moreover, as explained in [13] the same argument works for morphisms $\mathsf{Sub}(N \Rightarrow N, R) \to N$ for any *effectively separable* subset R, and this idea provides an important jump-off point for our approach. In Section 4.1, we introduce the general notion of a Δ-*separable* subset of $\mathbb{N}^{\mathbb{N}}$; these include typical 'tame' subsets that arise in mathematical practice, and we illustrate the scope of the concept by numerous examples. In Section 4.2, we shall show that under a combination of axioms that holds in almost all models of interest (and in particular without Normalizability), a KLS-style argument can be used to show regularity at $\mathbb{N}, R$ for all Δ-separable $R \subseteq \Delta$. In Section 4.3, we will show that under certain axioms including T-Normalizability, another variation on the KLS argument yields regularity at $\mathbb{N}, R$ for *all* $R \subseteq \Delta$. In Section 4.4, we will show that in the presence of P-Normalizability, much simpler arguments (modelled on Myhill-Shepherdson rather than KLS) can be used to show this much and a little more.

Let us comment briefly on how this relates to what can be done in the presence of simulations as in Section 2.7. If we knew, for example, that A admitted a well-behaved simulation in K_1, then the required upper bound on the 2-spectrum would be almost immediate from the classical KLS theorem. So in one sense, we are simply substituting a general axiomatic proof of KLS for the usual one; and although the theorems we obtain are new *as general results*, they do not really yield anything new in typical concrete instances. The stronger results possible in the presence of T-Normalizability appear to improve on classical KLS by dispensing with the separability requirement, but again this corresponds in concrete cases only to a mild and easy extension of existing results.

In the continuous case, where A admits a well-behaved simulation in K_2 or some submodel thereof, the upper bound on the 2-spectrum can be inferred even more straightforwardly from continuity in K_2. So in this case, we are having to work harder in our axiomatic setting to make up for the absence of a simulation.

In summary, then, our present results at second-order types do not in themselves yield much actual new information on particular models. Rather, the significance of our theorems is conceptual and methodological: they show that the axioms identified in Section 2 form a suitable basis for the development of the relevant theory at a general level. The extensions to higher types to be presented in Part II will provide further support for this claim, as well as yielding some genuinely new contributions to computability theory in particular cases.

4.1 Δ-separable sets

We start by introducing the concept of a Δ-separable subset of $\mathbb{N}^{\mathbb{N}}$, generalizing a familiar concept of effectively separable subset. Informally, a set R will be considered separable if the set of 1-codes a that match some $f \in R$ is enumerable, and moreover some such $f \in R$ can be computed from a:

Definition 33 *Suppose $\Delta \subseteq \mathbb{N}^{\mathbb{N}}$ is any set closed under Turing computation. A subset $R \subseteq \mathbb{N}^{\mathbb{N}}$ is* Δ-separable *if there exist a Δ-decidable set $H \subseteq \mathbb{N}$ and a uniformly A-computable family of functions $\zeta_0, \zeta_1, \ldots \in \Delta$ such that:*

1. *if $i \in H$ then $\zeta_i \in R$ (it will not matter what ζ_i is when $i \notin H$),*

2. for all $f \in R$ and $k \geq 0$, there exists $i \in H$ such that $\zeta_i(j) = f(j)$ for all $j < k$.

In this situation, we call the family $(\zeta_i \mid i \in H)$ a Δ-enumerable basis *for R.*

Once again, the set H is needed only to cover the case $R = \emptyset$ in a uniform and constructive way. We may also write $\widehat{\zeta_i}$ for some fixed choice of realizer for ζ_i computable from $\widehat{i}$ within A.

We say R is *effectively separable* if it is $\mathbb{N}^{\mathbb{N}}_{\mathrm{eff}}$-separable; clearly, any such R is Δ-separable for any Δ. Note too that in the above setting each ζ_i is in Δ, so if R is Δ-separable then so is $R \cap \Delta$ with the same enumerable basis.

We now give some examples to show the scope of this concept:

Examples 34 (i) If $\Delta = \mathbb{N}^{\mathbb{N}}$, any $R \subseteq \Delta$ at all is trivially Δ-separable: take an offset enumeration of all 1-codes a such that $R \cap U_a$ is inhabited, then use the axiom of choice to pick some $\zeta_i \in R \cap U_a$ for all such a. By assembling these into a single function $\zeta \in \Delta$, we see that ζ_i is A-computable uniformly in i. So in this case, the analysis below will apply to the whole of the 2-spectrum below $\mathbb{N}$.

(ii) Let us say a function $f : \mathbb{N} \to \mathbb{N}$ is *eventually zero* if there is some n such that $f(i) = 0$ for all $i \geq n$. It is easy to define a uniformly primitive recursive enumeration $\zeta_0, \zeta_1, \ldots$ of the eventually zero functions such that whenever i is a 1-code, ζ_i matches i; we shall call this the *standard basis*. Clearly this serves as an effectively enumerable basis for $R = \mathbb{N}^{\mathbb{N}}$, or indeed for any $R \subseteq \mathbb{N}^{\mathbb{N}}$ whatsoever that contains all the ζ_i.

(iii) By a Δ-*open* set $R \subseteq \mathbb{N}^{\mathbb{N}}$, we shall mean one such that there exists $e \in \Delta$ with $R = \bigcup\{U_{e(i)} \mid e(i) \text{ a 1-code}\}$. Then any Δ-open set is Δ-separable: indeed, it is easy to construct an enumerable base for each $U_{e(i)}$ uniformly in i, and these may be readily combined to yield an enumerable base for R.

(iv) Any Δ-enumerable (and hence countable) subset of Δ is Δ-separable: in this case, we may take the set itself as its own basis.

(v) Suppose we use functions $\mathbb{N} \to \mathbb{N}$ to represent real numbers in some standard way (e.g. via Cauchy sequences of rationals with fixed rate of convergence). Then all sets of the following kinds are effectively separable (we leave the construction of suitable bases as an exercise): enumerable unions of open intervals; enumerable unions of closed intervals; the set of rationals; the set of irrationals.

The following proposition affords another rich source of examples of separable sets; it provides our first example of a typical use of Continuity, and will itself be used in Section 4.3 below. We shall say a partial function $F : \Delta \rightharpoonup \mathbb{N}$ is *strictly represented* by $\widehat{F} \in A(\overline{2})$ if whenever $\widehat{f} \in N^N$ represents $f \in \Delta$, we have $\widehat{F} \cdot \widehat{f} = \widehat{m}$ iff $F(f) = m$. (Thus, if $F(f) \uparrow$, then $\widehat{F} \cdot \widehat{f} \Uparrow$.)

Proposition 35 *Suppose A satisfies Continuity, Enumeration and Restriction, and suppose $F : \Delta \rightharpoonup \mathbb{N}$ is strictly represented by some $\widehat{F} \in A(\overline{2})$. Then* $\mathrm{dom}\, F \subseteq \Delta$ *is Δ-separable.*

Proof Let $R = \mathrm{dom}\, F$, let $\widehat{F}$ strictly represent F, and let $\top, \bot$ be as in Section 2.2. We first use Continuity to show that if a is a 1-code and $R \cap U_a$ is inhabited then it is inhabited by an eventually zero function. Suppose $f \in R \subseteq U_a$ with $F(f) = n$, and for any $\alpha \in \top \in \bot$ define

$$f_\alpha(i) \;=\; \begin{cases} f(i) & \text{if } i \in \mathrm{dom}\, a \text{ and } i < t(\alpha) \\ 0 & \text{otherwise} \end{cases}$$

Note that the condition $i < t(\alpha)$ can be tested within A by checking whether $\alpha(\widehat{k}) = \widehat{0}$ for all $k \leq i$. Thus f_α is realizable in A by some $\widehat{f_\alpha}$ constructed uniformly in α. Moreover, if $\alpha = \bot$ then $f_\alpha = f$, so $\widehat{F} \cdot \widehat{f_\alpha} = \widehat{F(f)} = \widehat{n}$. Hence by Continuity there is some $\alpha \in \top$ with $\widehat{F} \cdot \widehat{f_\alpha} = \widehat{n}$, which means that $f_\alpha \in \mathrm{dom}\, F$. Finally, it is clear by construction that $f_\alpha \in U_a$, so $f_\alpha \in R \cap U_a$ and also that f_α is eventually zero.

Now suppose $i \mapsto \zeta_i$ is the standard enumeration of eventually zero functions, and form within $A(\overline{1})$ the operation $\lambda i.F(\zeta_i)$. Applying Enumeration to this, we obtain an offset enumeration in Δ of all pairs $\langle i, m_i \rangle$ where $\zeta_i \in R$ and $F(\zeta_i) = m_i$, and it is now straightforward to collect all such ζ_i into a Δ-enumerable basis for R. □

Under mild hypotheses, any Δ-open subset of Δ as in Example 34(iii) arises as the domain of some such F. However, there may also be more exotic domains that arise in this way: in the case $A = K_1$, there is an example due to Friedberg [9] of a partial function F as above whose domain is not open.

4.2 The non-normalizable case

We now consider what may be said about the 2-spectrum below $\mathbb{N}$ in the absence of a Normalizability Axiom; the results here will therefore apply even to extremely intensional models such as K_1 or models for PCF+timeout. We will see that a direct adaptation of the classical KLS proof allows us to get a certain distance using no axioms except Continuity. We then appeal to the Collection Axiom to complete the proof that certain parts of the 2-spectrum are indeed regular.

The original KLS theorem concerned, in effect, the 2-spectrum of K_1 at $\mathbb{N}, R$ for an arbitrary effectively separable set $R \subseteq \mathbb{N}^{\mathbb{N}}_{\mathrm{eff}}$.[5] We here adapt this proof to the general setting of a model A and a Δ_A-separable subset $R \subseteq \Delta_A$.

Suppose first that A is any model satisfying the Continuity Axiom, with $\top, \bot \subseteq A(\overline{1})$ and $t : \top \cup \bot \to \mathbb{N}_\infty$ as in Section 2.2. We write Δ for Δ_A.

The following shows how part of the classical KLS argument goes through in this setting.

Theorem 36 *Assume A satisfies Continuity. Suppose $R \subseteq \Delta$ is Δ-separable, and $F \in \mathsf{S}^2(\mathbb{N}, R)$. Then F is continuous on R: for any $f \in R$, there exists a 1-code a matching f such that $F(f') = F(f)$ for all $f' \in U_a \cap R$.*

Proof Suppose $R \subseteq \mathsf{S}^1(\mathbb{N})$ is Δ-separable, with $(\zeta_i \mid i \in H)$ an enumerable basis for R. Suppose also that $F \in \mathsf{S}^2(\mathbb{N}, R)$, and take $\widehat{F} \in A(\overline{2})$ a realizer for F.

By an *initial 1-code a*, we shall mean one of the form $\langle \langle 0, p_0 \rangle, \ldots, \langle r-1, p_{r-1} \rangle \rangle$. We write $r_a = r$ for the length of this 1-code, and (as usual) U_a for the associated neighbourhood

[5] The original treatment in [13] was at essentially the level of generality considered here; in [16] the result was stated only for a smaller class of effectively separable sets. We also point out a minor error in the statement of [16, Corollary 9.2.6]: the functional F' should be declared to be partial, i.e. of type $\mathbb{N}^{\mathbb{N}} \rightharpoonup \mathbb{N}$.

$\{f \in \mathbb{N}^{\mathbb{N}} \mid a \text{ matches } f\}$. Given $m \in \mathbb{N}$, we shall say $\langle a, m\rangle$ is a *suitable graph element* for F if for any $f \in R \cap U_a$ we have $F(f) = m$.

We will describe a test performable within A that can be used to recognize suitable graph elements. For this purpose, we shall represent our initial 1-codes in a rather roundabout way, by means of pairs (f, α) where $f \in \Delta$ and $\alpha \in \top$. To any such pair we associate the initial 1-code $a(f, \alpha) = \langle\langle 0, f(0)\rangle, \ldots, \langle t(\alpha) - 1, f(t(\alpha) - 1)\rangle\rangle$. Clearly, any initial 1-code is associated with infinitely many (f, α).

Given any $\widehat{f} \in N^N$ realizing $f \in \Delta$, $\alpha \in \top \cup \bot$ and $m \in \mathbb{N}$, we may construct an element $\widehat{\xi}_{\widehat{f},\alpha,m} \in A(\overline{1})$ that behaves as follows when applied to $\widehat{j} \in N$:

- If $j < t(\alpha)$ (that is, $\alpha(\widehat{k}) = \widehat{0}$ for all $k \leq j$), return $\widehat{f(j)}$.
- Otherwise, search for an $i \in H$ such that $\zeta_i(k) = f(k)$ for all $k < t(\alpha)$ and $F(\zeta_i) \neq m$. (Recall that for all $i \in H$ we have $\zeta_i \in R$, so that $\widehat{F} \cdot \widehat{\zeta_i} \Downarrow$.)
- If such an i is found, return $\widehat{\zeta_i(j)}$.

Notice that '$\zeta_i(k) = f(k)$ for all $k < t(\alpha)$' is equivalent to '$\zeta_i \in U_{a(f,\alpha)}$'. Notice too that exactly the same search for i will be performed for any $j \geq t(\alpha)$, so that if this search succeeds, the same i will be found irrespective of j.

Clearly $\widehat{\xi}_{\widehat{f},\alpha,m}$ may be computed within A uniformly in $\widehat{f}, \alpha, \widehat{m}$.

We shall now say that a triple $(\widehat{f}, \alpha, m)$ *passes the test* if $\alpha \in \top$ and $\widehat{F} \cdot \widehat{\xi}_{\widehat{f},\alpha,m} = \widehat{m}$. We will show that any triple passing the test yields a suitable graph element for F, and moreover that enough graph elements arise in this way to form a complete set graph of F.

Claim 1: If $(\widehat{f}, \alpha, m)$ passes the test where $\widehat{f}$ realizes f, then $\langle a, m\rangle$ is a suitable graph element, where $a = a(f, \alpha)$.

Proof: We wish to show that for any $f' \in U_a \cap R$ we have $F(f') = m$. We first show this for the case where f' is one of our functions ζ_i for $i \in H$.

Suppose $\zeta_i \in U_a \cap R$ for $i \in H$, so that $F(\zeta_i) \in \mathbb{N}$, and let us suppose for contradiction that $F(\zeta_i) = m' \neq m$. Then since also $i \in H$ and $\zeta_i \in U_{a(f,\alpha)}$ and $i \in H$, the search in the definition of $\widehat{\xi}_{\widehat{f},\alpha,m}$ will conclude at i or earlier. Let i' be the index returned by this search, so that $\zeta_{i'} \in U_a$ and $F(\zeta_{i'}) \neq m$. It is easy to see that $\widehat{\xi}'_{\widehat{f},\alpha,m}$ represents $\zeta_{i'}$ where $\zeta_{i'} \in R$; hence $\widehat{F} \cdot \widehat{\xi}_{\widehat{f},\alpha,m} = \widehat{F(\zeta_{i'})} \neq \widehat{m}$, contradicting that (f, α, m) passes the test. Thus $F(\zeta_i) = m$ after all.

Next, consider an arbitrary $f' \in U_a \cap R$, so that $F(f') \in \mathbb{N}$, and again suppose for contradiction that $F(f') = m' \neq m$. We shall use Continuity to obtain some $\zeta_i \in U_a \cap R$ close to f' with $F(\zeta_i) = m'$, contrary to what we have shown above. For each $\beta \in \top \cup \bot$, define $f'_\beta \in A(\overline{1})$ by

$$f'_\beta(n) \;=\; \begin{cases} f'(n) & \text{if } n < r_a \text{ or } n < t(\beta) \\ \zeta_i(n) & \text{otherwise, where } h(i) = a(f', \beta') + 1 \\ & \quad\text{with } t(\beta') = \max(r_a, t(\beta)) \end{cases}$$

Note again that $n < t(\beta)$ is equivalent to $\beta(\widehat{k}) = \widehat{0}$ for all $k \leq n$, and also that in the second case above, an i with the required properties will exist and can be found by a search through

h. We can therefore compute a realizer $\widehat{f'}_\beta$ for f'_β within A uniformly in β and any realizer for f'. Moreover, when $\beta \in \bot$, we have $t(\beta) = \infty$ so $f'_\beta = f'$, so that $\widehat{F} \cdot \widehat{f'}_\beta = \widehat{F(f')} = \widehat{m'}$. Hence by the Continuity Axiom, there must exist $\beta \in \top$ with $\widehat{F} \cdot \widehat{f'}_\beta = \widehat{m'}$. But now f'_β is equal to $\zeta_i \in R$ for a certain $i \in H$; thus $F(\zeta_i) = F(f'_\beta) = m'$. Also $\zeta_i \in U_a$, since $f'_\beta \in U_a$ by construction. But this is a contradiction, since we have shown above that $F(\zeta_i) = m$ for all $\zeta_i \in U_a \cap R$. Thus $F(f') = m$ after all. This completes the proof of Claim 1.

Claim 2: For any $f \in R$ with $F(f) = m$ and any $\widehat{f}$ realizing f, there exists $\alpha \in \top$ such that $(\widehat{f}, \alpha, m)$ passes the test and $f \in U_{a(f,\alpha)}$.

Proof: It is easy to see that if $\alpha \in \bot$ then $\widehat{\xi}_{\widehat{f},\alpha,m}$ represents f, so that $\widehat{F} \cdot \widehat{\xi}_{\widehat{f},\alpha,m} = \widehat{m}$. Hence by Continuity there exists $\alpha \in \top$ with $\widehat{F} \cdot \widehat{\xi}_{\widehat{f},\alpha,m} = \widehat{m}$, so that $(\widehat{f}, \alpha, m)$ passes the test. That $f \in U_{a(f,\alpha)}$ is automatic if $f \in \Delta$ and $\alpha \in \top$ by definition of $a(f, \alpha)$. This proves Claim 2.

The theorem itself now follows easily. Given $f \in R$, take any $\widehat{f}$ realizing f, set $m = F(f)$, then use Claim 2 to pick α such that $(\widehat{f}, \alpha, m)$ passes the test and $f \in U_a$ where $a = a(f, \alpha)$. Claim 1 now says that $\langle a, m \rangle$ is a suitable graph element, which is to say that $F(f') = F(f)$ for all $f' \in U_a \cap R$. $\Box$

Some of the constructions in the above proof may be usefully generalized a little. If $\widehat{f} \in N^N$ is a realizer for f, let us allow the notation $a(\widehat{f}, \alpha)$ in place of $a(f, \alpha)$, so that

$$a(\widehat{f}, \alpha) \;=\; \langle \langle 0, c_0 \rangle, \ldots, \langle r-1, c_{r-1} \rangle \rangle \;\text{ iff }\; r = t(\alpha) \text{ and } \forall k < r.\; \widehat{f} \cdot \widehat{k} = \widehat{c_k} \;.$$

We can use this to also serve as a definition of $a(\widehat{f}, \alpha)$ even for $\widehat{f} \notin N^N$, provided that $\alpha \in \top$ and $\widehat{f} \cdot \widehat{k} \Downarrow$ for all $k < t(\alpha)$: it does not matter how $\widehat{f}$ behaves on larger $\widehat{k}$. Let us take

$$P \;=\; \{ (\widehat{f}, \alpha) \;\mid\; \alpha \in \top,\; \widehat{f} \cdot \widehat{k} \Downarrow \text{ for all } k < t(\alpha) \} \;.$$

We now observe that the definition of $\widehat{\xi}_{\widehat{f},\alpha,m}$ makes good sense as long as $(\widehat{f}, \alpha) \in P$, and indeed that the proof of Claim 1 also goes through under this weaker hypothesis. There is thus no danger of generating unsuitable graph elements if we allow these additional elements $\widehat{f}$, provided the condition $(\widehat{f}, \alpha) \in P$ is respected.

This suggests the idea of trying to generate a set graph for F as the proper range of an operation within A mapping each suitable $(\widehat{f}, \alpha, m)$ to $\langle a(\widehat{f}, \alpha), m \rangle$. To make this work, we now adopt the Iterated Restriction Axiom from Section 2.6.

Lemma 37 *Assume A satisfies Continuity and Iterated Restriction, and suppose R and F are as in Theorem 36. Then there exists $\Psi \in A(\overline{1} \to \overline{1} \to \overline{0} \to \overline{0})$ whose proper range*

$$\widehat{G} \;=\; \{ \Psi \cdot x \cdot y \cdot z \;\mid\; x : \overline{1},\; y : \overline{1},\; z : \overline{0},\; \Psi \cdot x \cdot y \cdot z \Downarrow \}$$

exactly corresponds to some set graph $G \subseteq \mathbb{N}$ for F.

Proof Let $\widehat{a} \in A(\overline{1} \to \overline{1} \to \overline{0})$ be a realizer for a as a function $P \to N$, constructed in the obvious way using iteration. Now consider the element

$$\Phi \;=\; \lambda^* xyz.\; pair\; (\widehat{a}\; x\; y)\; z \;\in\; A(\overline{1} \to \overline{1} \to \overline{0} \to \overline{0}) \;.$$

From the discussion above, it is clear that $\Phi \cdot x \cdot y \cdot z$ will return the numeral for a suitable graph element for F, provided all of the following conditions are satisfied:

1. $z \in N$ (say $z = \widehat{m}$).

2. $y \in \top$.

3. $(x, y) \in P$.

4. $\widehat{F} \cdot \widehat{\xi}_{x,y,m} \Downarrow$ and $\widehat{F} \cdot \widehat{\xi}_{x,y,m} = \widehat{m}$.

We have also seen that enough graph elements can be thus obtained to cover every $f \in R$.

We claim that each of the above conditions can be checked via the convergence of some expression involving x, y, z, assuming that all the previous conditions are satisfied. Condition 1 is readily checked via a convergence test. For Condition 2, we first check that the value for $t(\alpha)$ computed by minimization converges, say to $\widehat{u}$, and then used Iterated Restriction to check that indeed $\alpha(\widehat{u}) = \widehat{1}$ and $\alpha(k) = \widehat{0}$ for all $k < u$. Having obtained this numeral $\widehat{u}$, we may now again use Iterated Restriction to test Condition 3 by checking that $\widehat{f} \cdot \widehat{k} \Downarrow$ for all $k < u$. Finally, Condition 4 is easily checked via two simple convergence tests.

Putting all this together, we may construct an expression $e(x, y, z)$ that converges iff all four conditions are satisfied. Finally, we may define

$$\Psi \ = \ \lambda^* xyz.\ (\Phi xyz) \upharpoonright e(x, y, z)$$

and it is clear that the proper range of Ψ yields a graph for F. $\square$

We have thus obtained a set graph for F. With our current axioms, we cannot collect the graph elements into an enumerated graph for F, as we have no way within A to enumerate all, or sufficiently many, elements $\widehat{f} \in A(\overline{1})$. (Even if we can arrange that an enumerable family of functions $f \in \Delta$ suffices, as we shall in the proof of Theorem 42, it seems that we cannot enumerate sufficiently many realizers for these functions.) However, it is at this point that the Collection Axiom exactly meets our need:

Theorem 38 *Assume A satisfies Continuity, Iterated Restriction, Collection and Enumeration, and suppose $R \subseteq \Delta$ is Δ-separable. Then A is 2-regular at $\mathbb{N}, R$.*

Proof Let F be any function $R \to \mathbb{N}$. We have seen in Proposition 30 that if F has a graph then $F \in \mathsf{S}^2(\mathbb{N}, R)$. Conversely, if $F \in \mathsf{S}^2(\mathbb{N}, R)$, then by Lemma 37 there exists $\Psi : \overline{1} \to \overline{1} \to \overline{0} \to \overline{1}$ whose proper range yields a set graph G for F. Now apply the $\overline{1}, \overline{1}, \overline{0}$-Collection Axiom to this Φ to obtain $f \in A(\overline{1})$ whose proper range gives a graph of F; then apply Enumeration to obtain $g \in N^N$ whose offset range is a graph of F. Finally, since this graph is certainly inhabited, we may (within A) remove the offsetting to obtain an ordinary enumerated graph $g' \in \Delta$ for F. $\square$

Corollary 39 *If A satisfies the above axioms and $\Delta_A = \mathbb{N}^{\mathbb{N}}$, then A is 2-regular below $\mathbb{N}$.*

We emphasize again that the combination of the above four axioms holds in almost all our models of interest, with the exception of certain untyped λ-algebras in which Iterated Restriction is problematic, and with the qualification that the treatment of certain 'non-deterministic' models requires a few additional refinements as described in Section 6 below.

Remark 40 Our argument above is in essence nothing more than the classical KLS proof transposed to our setting. Indeed, our version can be seen as offering a conceptual analysis of the classical proof in terms of three ingredients: the proof that K_1 satisfies the Continuity Axiom (using the undecidability of the halting problem), the proof that certifiable local moduli exist and give rise to graph elements (which requires only the Continuity Axiom), and the construction for collecting these graph elements into an enumerable family (which exploits the specific fact that 'type 2' objects in K_1 are just natural numbers).

Since we have noted in Example 34(iv) that the class of Δ-separable sets includes the domains of all representable partial functions on Δ, it might be expected that Theorem 38 suffices to determine the entire 2-spectrum below $\mathbb{N}$, since (one might suppose) any partial function F whatever is present in $\mathsf{S}^2(\mathbb{N}, \mathrm{dom}\ F)$ and hence has a graph, which will then also be applicable to the restriction F to other sets $R \subseteq \mathrm{dom}\ F$. However, this is not the case: Proposition 35 applies only to operations $\widehat{F}$ whose domains are *extensional*, i.e. such that if $\widehat{f}, \widehat{f'} \in N^N$ realize the same $f \in \Delta$, then $\widehat{F} \cdot \widehat{f} \Downarrow$ iff $\widehat{F} \cdot \widehat{f'} \Downarrow$. There are many other elements $\widehat{F}$ whose domains appear 'ragged at the edges' if we try to view them as realizing partial functions on $\mathbb{N}^{\mathbb{N}}$, and such elements can contribute to the 2-spectrum at various non-separable R.

The following counterexample (inspired by Friedberg's) shows that under the present hypotheses, one cannot hope to extend Theorem 38 to arbitrary subsets of Δ.

Counterexample 41 Consider Kleene's first model K_1, the untyped PCA consisting of the natural numbers with (partial) Kleene application $\bullet$. Let us say a function $f : \mathbb{N} \to \mathbb{N}$ *attains its maximum at* i if $f(i) \geq f(j)$ for all $j \in \mathbb{N}$, and (when f is computable) write e_f for the smallest e such that $e \bullet j = f(j)$ for all j. Now consider the effectively non separable set

$$R_{\max} \;=\; \{f \in \mathbb{N}^{\mathbb{N}}_{\mathrm{eff}} \mid f \text{ attains its maximum at some } i \leq e_f\,\}$$

and define $F_{\max} : R_{\max} \to \mathbb{N}$ by taking $F(f)$ to be the maximum value $f(j)$ attained by f. Then $F_{\max}$ is realizable in K_1: given any realizer e for any $f \in R_{\max}$, the required value of $F_{\max}(f)$ is just $\max(e \bullet 0, e \bullet 1, \ldots, e \bullet e)$. However, it is an easy exercise to show that $F_{\max}$ cannot have a graph. Thus the spectrum of K_1 is not regular at $\mathbb{N}, R_{\max}$.

This marks a limit to the extent of the regularity phenomenon in the absence of Normalizability. We will see in Section 4.3 that T-Normalizability gives rise to regularity at all $\mathbb{N}, R$.

4.3 The T-normalizable case

For this subsection, we assume that A satisfies T-Normalizability as well as Continuity and Enumeration. Our main result here is that under these hypothesis, regularity extends to the whole of the 2-spectrum below $\mathbb{N}$. Our proof will appeal to Theorem 36, but in addition, T-Normalizability will play two distinct roles. Firstly, by working with normalized realizers, it becomes trivial that any $\widehat{F} \in A(\overline{2})$ behaves extensionally on type 1 realizers whenever it is defined, even outside the originally intended domain of F. Secondly, it allows us to work with a more minimalist version of the sets $\top, \bot$, with the property that $\top$ is enumerable within A; this will mean that Enumeration alone suffices to collect sufficiently many graph elements to form an enumerated graph of our functional F.

For the purpose of this section, it will be sufficient to work with the standard basis $\zeta_0, \zeta_1, \ldots$ consisting of all eventually zero functions. As before, we write $\widehat{\zeta_i} \in N^N$ for a realizer for ζ_i computed uniformly in i within A. We may also assume by construction that each $\widehat{\zeta_i}$ is normalized (that is, $norm \cdot \widehat{\zeta_i} = \widehat{\zeta_i}$ where $norm$ is as in Definition 8).

It will also be convenient to reformulate the Continuity Axiom slightly in this setting. For each $n \in \mathbb{N}_\infty$, let us take $\alpha_n \in A(\overline{1})$ to be the unique normalized element such that

$$\alpha_n \cdot \widehat{i} \;=\; \begin{cases} \widehat{0} & \text{if } i < n, \\ \widehat{1} & \text{if } i \geq n. \end{cases}$$

As an alternative to the usual sets $\bot, \top$ from Section 2.2, let us define

$$\bot_1 \;=\; \{\alpha_\infty\} \,, \qquad \top_1 \;=\; \{\alpha_n \mid n \in \mathbb{N}\} \,.$$

It is easy to see that there is an element $C \in A(\overline{1} \to \overline{1})$ mapping $\top$ to $\top_1$ and $\bot$ to $\bot_1$: take

$$C \;=\; \lambda^* \alpha.\, norm\, (\lambda^* i.\, \alpha(min\ j.\, j = i \vee \alpha(j) = \widehat{1})) \,.$$

The Continuity Axiom as stated in Definition 4 therefore immediately implies the following new version of Continuity: For any $F \in A(\overline{2})$, if $F \cdot \alpha_\infty = \widehat{m}$ then $F \cdot \alpha_n = \widehat{m}$ for some $n \in \mathbb{N}$. It is clear that the proof of Theorem 36 above goes through under this new form of Continuity, using $\top_1, \bot_1$ in place of $\top, \bot$.

Theorem 42 *Suppose A satisfies Continuity, T-Normalizability and Enumeration. Then A is 2-regular below $\mathbb{N}$.*

Proof Suppose $R \subseteq \mathsf{S}^1(\mathbb{N})$ and $F \in \mathsf{S}^2(\mathbb{N}, R)$; we wish to find a graph of F within $\Delta = \Delta_A$. Let $\widehat{F}' \in A((\mathtt{N} \to \mathtt{N}) \to \mathtt{N})$ be any realizer for F, and let $\widehat{F} = \lambda^* f.\, \widehat{F}'\,(norm\ f)$, so that $\widehat{F}$ also realizes F. Now let $R^+ \subseteq \mathsf{S}^1(\mathbb{N})$ be the set of f such that $\widehat{F} \cdot \widehat{f} \Downarrow$ for some (and hence any) $\widehat{f}$ realizing f, so that $R^+ \supseteq R$, and let $F^+ : R^+ \to \mathbb{N}$ be the extension of F that is realized by $\widehat{F}$. Since $\widehat{F}$ has normalization built in, it is clear that $\widehat{F}$ strongly represents F^+ in the sense of Proposition 35. We shall in fact construct a graph of F^+ in Δ; this will *a fortiori* be a graph for F, so we may henceforth forget about F.

By Proposition 35, R^+ is Δ-separable and indeed has a basis consisting of eventually zero functions. So by Theorem 36, we immediately have that F^+ is continuous on R^+. Furthermore, the proof of Theorem 36 (adapted to use $\top_1, \bot_1$) shows that the set

$$G \;=\; \{\ \langle a(f,\alpha), m\rangle \ \mid\ f \in R^+,\ \alpha \in \top_1,\ m \in \mathbb{N},\ \widehat{f} \text{ realizes } f,\ (\widehat{f}, \alpha, m) \text{ passes the test}\ \}$$

constitutes a set graph for F. As it stands, this set is not amenable to enumeration, as we have no way within A to enumerate a set of representatives for all $f \in R^+$. However, we will show that we still have enough graph elements if we restrict to the class of eventually zero f, which we can enumerate.

Claim: For any $f \in R^+$ with $F^+(f) = m$, there exist $\zeta_i \in R^+$ and $\alpha \in \top_1$ such that $(\widehat{\zeta_i}, \alpha, m)$ passes the test and $f \in U_{a(\zeta_i,\alpha)}$.

Proof: Given f, m as above, take $\alpha \in \top_1$ so that (f, α, m) passes the test and $f \in U_a$ where $a = a(f, \alpha)$. Exactly as in the proof of Theorem 36, we may construct elements f_β for $\beta \in \top_1 \cup \bot_1$ such that $f_\beta = f$ for all $\beta \in \bot_1$, f_β is an eventually zero in R^+ for $\beta \in \top_1$, and $f_\beta \in U_a$ for all β. Since (f, α, m) passes the test, by Continuity there is some $\beta \in \top_1$ such that (f_β, α, m) passes the test. Moreover, since $f_\beta \in U_a$, we have $a(f_\beta, \alpha) = a$. But $f \in U_a$, so $f \in U_{a(f_\beta, \alpha)}$. Finally, since f_β is eventually zero, we may pick i such that $\zeta_i = f_\beta \in R^+$, and this proves the claim.

Our final task is to collect together these elements to obtain an actual graph in Δ. Construct within A an enumeration $(i_0, \alpha_0, m_0), (i_1, \alpha_1, m_1), \ldots$ of all triples (i, α, m) with $\alpha \in \top_1 \cup \bot_1$ and ζ_i in our basis for R^+, and compose this with the operation

$$(i, \alpha, m) \;\mapsto\; \widehat{F} \cdot \widehat{\xi}(\zeta_i, \alpha, m) \; .$$

Write h for the resulting element of $A(\mathtt{N} \to \mathtt{N})$. Applying Enumeration to h, we obtain $g \in \Delta$ with offset range consisting of exactly the pairs $\langle j, m'_j \rangle$ where $\widehat{F} \cdot \widehat{\xi}(\zeta_{i_j}, \alpha_j, m_j)$ converges to $\widehat{m'_j}$. By further easy programming within A, we may transform this into $g' \in \Delta$ whose offset range is exactly the set of $\langle a(\zeta_i, \alpha), m \rangle$ such that (ζ_i, α, m) passes the test, and hence into $g'' \in \Delta$ whose ordinary range is the set of such $\langle a(\zeta_i, \alpha), m \rangle$, perhaps with $\langle\rangle$ added.

It is now easy to see that $g'' \in \Delta$ serves as a graph of F. Condition 1 of Definition 29 follows immediately from the Claim above; Condition 2 holds because G is a set graph for F^+; and Condition 3 is vacuous here since $Q = \mathbb{N}$. $\square$

4.4 The P-normalizable case

P-Normalizability gives us a very good grasp indeed of elements of type $\overline{1}$: since *every* $f \in A(\overline{1})$ represents some partial function $\mathbb{N} \rightharpoonup \mathbb{N}$, the normalizer *norm* in effect collapses the whole of $A(\overline{1})$ into a set of such partial functions. Moreover, in the presence of Continuity and Restriction, it is not hard to show that any A-representable function $(\mathbb{N} \rightharpoonup \mathbb{N})$ enjoys certain monotonicity and continuity properties.

If $f \in A(\overline{1})$, we shall write $\check{f} : \mathbb{N} \rightharpoonup \mathbb{N}$ for the associated partial function: $\check{f}(n) = m$ iff $f(\widehat{n}) = \widehat{m}$.

Lemma 43 *Assume A satisfies Continuity, Restriction and P-Normalizability, and suppose $F \in A(\overline{2})$.*

(i) If $f \in A(\overline{1})$ is normalized and $F \cdot f = \widehat{n}$, then there exists $d \in A(\overline{1})$ such that $\check{d}$ is finite, $\check{f}$ extends $\check{d}$ and $F \cdot d = \widehat{n}$.

(ii) If $d \in A(\overline{1})$ with $\check{d}$ finite and $F \cdot d = \widehat{n}$, then for any normalized $f \in A(\overline{1})$ with $\check{f}$ extending $\check{d}$, we have $F \cdot f = \widehat{n}$.

Proof (i) Using Restriction as indicated in Section 2.6, define $K \in A(\overline{1} \to \overline{1})$ so that $K \cdot f \downarrow$ for all $f \in A(\overline{1})$, $K \cdot f \cdot \widehat{j} = \widehat{0}$ if $f \cdot \widehat{j} = \widehat{0}$, and $K \cdot f \cdot \widehat{j} \Uparrow$ otherwise. Now given $f \in A(\overline{1})$, define $E_f \in A(\overline{1} \to \overline{1})$ by

$$E_f \;=\; \lambda^*\alpha.\; norm\,(\lambda^* j.\, (fj) \upharpoonright (K\alpha j)) \; .$$

If $\alpha \in \bot$ then $K \cdot \alpha \cdot \widehat{j} = \widehat{0}$ for all j, so $E_f \cdot \alpha$ represents $\check{f}$, whence $E_f = norm \cdot f$.

Now suppose $F \cdot f = \widehat{n}$. Then $F \cdot (E_f \cdot \alpha) = \widehat{n}$ for all $\alpha \in \bot$, so by continuity we may take $\alpha \in \top$ such that $F \cdot (E_f \cdot \alpha) = \widehat{n}$. But now $K \cdot \alpha \cdot \widehat{j} \Downarrow$ iff $j < t$, where $t = t(\alpha) < \infty$; hence $d = E_f \cdot \alpha$ represents a finite function $\check{d} : \mathbb{N} \rightharpoonup \mathbb{N}$, where $\check{d}(j) \simeq f(j)$ for $j < t$, and $\check{d}(j) \uparrow$ for $j \geq t$. Thus $\check{f}$ extends $\check{d}$, and we have seen that $F \cdot d = \widehat{n}$.

(ii) Suppose $\check{d}$ is finite and $F \cdot d = \widehat{n}$, and consider a normalized f with $\check{f}$ extending $\check{d}$. We may assume d is normalized. Write $\{q_0, \ldots, q_{r-1}\}$ for the domain of d, set $p_i = d(q_i)$ for each i, and consider $G : \overline{1} \to \overline{1}$ defined by

$$G \;=\; \lambda^* h.\ \mathit{norm}\,(\lambda j.\ \mathit{if}\ (\mathit{eq}\ j\ \widehat{q_0})\ \mathit{then}\ \widehat{p_0}\ \mathit{else}\ \cdots \mathit{if}\ (\mathit{eq}\ j\ \widehat{q_{r-1}})\ \mathit{then}\ \widehat{p_{r-1}}\ \mathit{else}\ (f j) \upharpoonright (h\,\widehat{0}))\ .$$

Clearly, if $h \in \bot_0$ then $G \cdot h$ is a normalized realizer for $\check{d}$, so $G \cdot h = d$. On the other hand, if $h \in \top_0$ then $G \cdot h$ is a normalized realizer for $\check{f}$, so $G \cdot h = f$. Thus $d \ll f$ in the sense of Lemma 15, and since $F \cdot d = \widehat{n}$ it follows that $F \cdot f = \widehat{n}$. $\Box$

We may now obtain a result for representable *partial* functionals which has no counterpart in the non-P-normalizable setting. Recall that Definition 29 defined a notion of graph for partial functionals. We shall write $\Delta_\bot$ for the set of functions $\mathbb{N} \to \mathbb{N}_\bot$ representable in A; we will henceforth feel free to identify normalized elements of $A(\overline{1})$ with elements of $\Delta_\bot$.

Theorem 44 (Myhill-Shepherdson property) *Assume A satisfies Continuity, Restriction, P-Normalizability and Enumeration. Suppose $F : \Delta_\bot \rightharpoonup \mathbb{N}$ has a realizer $\widehat{F} \in A(\overline{2})$, in the sense that for all $f \in \Delta_\bot$, $\widehat{F} \cdot f = \widehat{m}$ iff $F(f) = m$. Then F has a graph in Δ.*

Proof Suppose $F, \widehat{F}$ are as above. Take $\gamma \in A(\overline{1})$ an enumeration of all 1-codes for finite partial functions $\mathbb{N} \to \mathbb{N}$, and take $\delta \in A(\overline{0} \to \overline{1})$ mapping any such 1-code to a realizer for the finite partial function it codes. Now consider $\epsilon \in A(\overline{1})$ defined by

$$\epsilon \;=\; \lambda i.\ \langle \gamma i, \widehat{F}(\delta i) \rangle \upharpoonright (\widehat{F}(\delta i))\ .$$

Then the proper range of ϵ consists of all graph elements $\langle a, m \rangle$ such that $F(d_a) = m$, where d_a denotes the finite partial function with 1-code a. By Theorem 43(ii), each of these is a suitable graph element for F, and by Theorem 43(i), every f such that $F(f)\downarrow$ is covered by some such graph element. Applying Enumeration to ϵ, we may obtain an actual graph $g \in \Delta$ for F. $\Box$

Recall that the classical Myhill-Shepherdson theorem also includes what is essentially the converse property: if F has a computable graph then F has a realizer. However, this fails in our general setting, since there are functionals with computable graphs that require 'parallel' operations for their evaluation.

We may now easily deduce 2-regularity below $\mathbb{N}$. We shall give a version of this that strengthens what we already know from Theorem 42:

Definition 45 *Suppose $R \subseteq \mathsf{S}^1(\mathbb{N})$.*

(i) A graph g for $F \in \mathsf{S}^2(\mathbb{N}, R)$ is a strong graph *if the following strengthening of Condition 2 from Definition 29 holds: For every $\langle a, m \rangle$ in the range of g and every $f \in \Delta$ matching a (without requiring $f \in R$), we have $F(f) = m$.*

(ii) The spectrum of A is strongly regular at $\mathbb{N}, R$ *if every $F \in \mathsf{S}^2(\mathbb{N}, R)$ has a strong graph.*

Theorem 46 *Assume A satisfies Continuity, Enumeration, Restriction and P-Normalizability. Then S^2_A is strongly regular below $\mathbb{N}$.*

PROOF Suppose $R \subseteq \mathsf{S}^1(\mathbb{N})$ and $F \in \mathsf{S}^2(\mathbb{N}, R)$, and take $\widehat{F} \in A(\overline{2})$ a realizer for F. Let $g \in \Delta$ be a graph of $\widehat{F}$ as in Theorem 44 (obtained without reference to R). It is easy to see that g serves as a graph for F: if $f \in R$ then $F(f) \downarrow$ so that $\widehat{F} \cdot f \Downarrow$, whence f is covered by an element $\langle a, m \rangle$ of g, and the suitability of such graph elements over the whole of $\Delta_\perp$ entails the strong graph property. □

5 The 2-spectrum below $Q \subseteq \mathbb{N}$

Thus far we have concentrated on the 2-spectrum below $\mathbb{N}$. We now consider what can be said about the 2-spectrum below an arbitrary $Q \subseteq \mathbb{N}$. If Q is inhabited and *enumerable* within A, say via a surjection $\rho : \mathbb{N} \to Q$ in Δ, the 2-spectrum below Q is easily obtained from that below $\mathbb{N}$. In this case, ρ has a one-sided inverse τ (with $\rho \circ \tau = id_Q$), computable in A by minimization, so that $\mathsf{Sub}(N, Q)$ is a retract of N in $\mathcal{M}od(A)$, whence also $\mathsf{R}^1(Q)$ is a retract of $\mathsf{R}^1(\mathbb{N}) = N \Rightarrow N$. For any $R \subseteq \mathsf{S}^1(Q)$, this restricts to a retraction $\mathsf{Sub}(\mathsf{R}^1(Q), R) \lhd \mathsf{Sub}(\mathsf{R}^1(N), R')$, where $R' = \{f \in \mathbb{N}^\mathbb{N} \mid f \upharpoonright_Q \in R\}$, and it follows that the elements of $\mathsf{S}^2(Q, R)$ are precisely those induced by elements of $\mathsf{S}^2(\mathbb{N}, R')$. Moreover, any graph $g \in \Delta$ for an element $F \in \mathsf{S}^2(\mathbb{N}, R')$ readily induces a graph $g' \in \Delta$ for the corresponding $F' \in \mathsf{S}^2(\mathbb{N}, R)$, simply by using the surjection ρ to translate graph elements as follows:

$$\langle\langle\langle q_0, p_0\rangle, \ldots, \langle q_{r-1}, p_{r-1}\rangle\rangle, m\rangle \quad \mapsto \quad \langle\langle\langle \rho(q_0), p_0\rangle, \ldots, \langle \rho(q_{r-1}), p_{r-1}\rangle\rangle, m\rangle$$

The use of ρ here ensures that Condition 3 of Definition 2-graph is satisfied. (The translation process might of course give rise to elements involving 'inconsistent' 1-codes: these may be either suppressed in g' or left as they are, since they do no particular harm.)

Let us now declare a set $R \subseteq \mathsf{S}^1(Q)$ to be Δ-*separable* if the induced set

$$R' \;=\; \{f \mid f \upharpoonright_Q \in R\} \;\subseteq\; \mathbb{N}^\mathbb{N}$$

is Δ-separable. From Theorem 38 we may now conclude:

Theorem 47 *Suppose A satisfies Continuity, Iterated Restriction, Collection and Enumeration. If $Q \subseteq \mathbb{N}$ is Δ_A-enumerable and $R \subseteq \mathsf{S}^1(Q)$ is Δ_A-separable, then S^2_A is regular at Q, R.*

Note that in the case $\Delta = \mathbb{N}^\mathbb{N}$, then (classically) *every* inhabited set $Q \subseteq \mathbb{N}$ is enumerable, and every $R' \subseteq \mathbb{N}^\mathbb{N}$ is separable. So we may conclude:

Theorem 48 *Under the hypotheses of Theorem 47, if $\Delta_A = \mathbb{N}^\mathbb{N}$, then A is 2-regular.*

We also have the counterpart to Theorem 47 in the T-Normalizable setting:

Theorem 49 *Suppose A satisfies Continuity, T-Normalizability and Enumeration. If $Q \subseteq \mathbb{N}$ is enumerable and $R \subseteq \mathsf{S}^1(Q)$ is arbitrary, then S^2_A is regular at Q, R.*

The situation when Q is not enumerable is harder to analyse. Although this perhaps takes us outside the realm of computable functionals arising in typical mathematical practice, having formulated our definition of the 2-spectrum it is natural to inquire how far the regularity phenomenon extends into it.

The following shows that even in some very well-behaved models, regularity can fail even at $Q, \mathsf{S}^1(Q)$ when Q is non-enumerable.

Counterexample 50 Work in the effective Scott model $\mathsf{PC}^{\mathrm{eff}}$, which satisfies P-Normalizability along with Continuity, Enumeration, Restriction and Collection. For our present purpose, let us say a statement is *provable* if some natural formalization of it is provable in (say) Peano Arithmetic. We also fix on some algorithmic procedure for searching for a proof of a given statement of PA, which will eventually find one if it exists.

Let Q be the complement of the halting set H; for convenience we suppose $0 \in Q$. Define $F : \mathsf{S}^1(Q) \to \mathbb{N}$ classically by

$$F(f) \;=\; \begin{cases} 0 & \text{if } f(0)+1 \in H \\ 1 & \text{if a proof of } f(0)+1 \in Q \text{ is found within time } f(f(0)+1) \\ 0 & \text{if } f(0)+1 \in Q \text{ but no proof of this is found within time } f(f(0)+1). \end{cases}$$

This is a valid definition in that it makes no reference to the behaviour of f outside its domain Q. Moreover, it is easy to see that a realizer for F within $\mathsf{PC}^{\mathrm{eff}}(\overline{2})$ is given by the following c.e. set of compact elements:

$$\begin{array}{l} \{((0 \mapsto q) \mapsto 0) \mid q+1 \in H\} \;\cup \\ \{((0 \mapsto q), (q+1 \mapsto t) \mapsto 1) \mid \text{`}q+1 \in Q\text{' proven within time } t\,\} \;\cup \\ \{((0 \mapsto q), (q+1 \mapsto t) \mapsto 0) \mid \text{`}q+1 \in Q\text{' not proven within time } t\,\} \end{array}$$

This is a consistent set of compacts because PA is sound. Thus $F \in \mathsf{S}^2(Q, \mathsf{S}^1(Q))$.

Suppose however that F had a computable graph g in the sense of Definition 29. Then for every $q \in \mathbb{N}$, g contains *either* an element $\langle\langle\langle 0, q\rangle, \ldots\rangle, 0\rangle$ where the '$\ldots$' does not involve a pair $\langle q+1, t\rangle$, *or* an element of the form $\langle\langle\langle 0, q\rangle, \langle q+1, t\rangle, \ldots\rangle, i\rangle$. (It may of course contain both.) Note that if $q+1 \in H$ then an element of the second kind cannot be present, and if provably $q+1 \in Q$ then an element of the first kind cannot be present. By searching for which kind of element appears first within g, we can thus obtain a decidable set D such that $H \subseteq D$ but $P \subseteq \overline{D}$, where P is the set of p for which $p \in Q$ is provable.

However, a Gödelian argument shows that no such decidable separation of H and P is possible. Clearly, the question is equivalent to whether there is a decidable set E of Σ^0_1 statements such that E contains all the true such statements and $\overline{E}$ all the provably false ones. (Recall that 'true' and 'provably true' coincide for Σ^0_1 statements.) Let $\phi_0, \phi_1, \ldots$ be some effective enumeration of all Σ^0_1 statements with a single free variable x, and let $\psi[q]$ be the statement '$\phi_q[x \mapsto q] \in \overline{E}$', framed as a Σ^0_1 statement. Take p so that $\phi_p = \psi[x]$; we now ask whether the Σ^0_1 statement $\psi[p]$ holds. If $\psi[p]$ is true then $\phi_p[x \mapsto p] \in \overline{E}$, i.e. $\psi[p] \in \overline{E}$, contradicting that E contains all true Σ^0_1 statements. On the other hand, if $\psi[p]$ is false then there is a finite computation showing that $\phi_p[x \mapsto p] \in E$, whence $\psi[p]$ is *provably* false; but then we have $\psi[p] \in E$, contradicting that $\overline{E}$ contains all provably false Σ^0_1 statements. So we have reached a contradiction.

We thus conclude that our functional F does not have a computable graph, and hence that $\mathsf{PC}^{\mathrm{eff}}$ is not regular at $\overline{H}, \mathsf{S}^1(\overline{H})$.

The above counterexample can be easily adapted to many other models that support interleaved or 'parallel' computation, such as K_1. Interestingly, however, the regularity phenomenon extends significantly further in the case of *sequential* computability models, to which we now turn our attention.

Definition 51 *(i) By a* sequential decision tree *we shall mean a countably branching tree in which*

- *each internal node is labelled with a* question $?q$ *where* $q \in \mathbb{N}$,
- *each leaf is labelled with an* answer $!m$ *where* $m \in \mathbb{N}$,
- *each edge is labelled with a number* p*, and distinct edges branching from the same internal node carry distinct labels.*

If T *is a sequential decision tree and* $f : \mathbb{N} \rightharpoonup \mathbb{N}$ *is a partial function, then* $T \bullet f$*, if defined, is the unique* m *such that there is a finite path* $?q_0, p_0, \ldots, ?q_{r-1}, p_{r-1}, !m$ *through* T *with* $f(q_i) = p_i$ *for each* i.

(ii) Suppose A *satisfies P-Normalizability and Restriction. We say* A *is* sequential *if for every* $F \in A(\overline{2})$*, the action of* F *on normalized elements of* $A(\overline{1})$ *agrees with that of some* Δ_A*-enumerable sequential decision tree* T_F.

It is easy to convince oneself that this definition covers the term models for intuitively 'sequential' languages such as PCF and its extensions with local state, continuations, coroutines etc., as well as the known game models for such languages. Another example is the model SR of *sequentially realizable* functionals (classically equivalent to the strongly stable functionals). The 'prototypical' sequential model is van Oosten's combinatory algebra B, which is explicitly built around the above notion of decision tree. Note in passing that any sequential model automatically satisfies Continuity.

Theorem 52 *If* A *is sequential, then* S^2_A *is regular at* $Q, \mathsf{S}^1(Q)$ *for any* $Q \subseteq \mathbb{N}$.

Proof Let $R = \mathsf{S}^1(Q)$, and suppose $F \in \mathsf{S}^2(Q, R)$. Take $\widehat{F} \in A(\overline{2})$ a realizer for F, and let $T_{\widehat{F}}$ be the corresponding Δ-enumerable sequential tree. We may assume without loss of generality that $T_{\widehat{F}}$ is irredundant, i.e. the same question is never asked twice on any path through F.

We claim that for every node $?q$ within $T_{\widehat{F}}$ we have $q \in Q$. Otherwise, take some node $?q$ of minimal depth within $T_{\widehat{F}}$ such that $q \notin Q$, and let $?q_0, p_0, \ldots, ?q_{r-1}, p_{r-1}$ be the path to this node. Define $f : \mathbb{N} \to \mathbb{N}_\perp$ by: $f(q_i) = p_i$ for each i; $f(q) = \perp$; and $f(j) = 0$ for all other j. Clearly f is realizable in $A(\overline{1})$. Moreover, f is total on Q, so we ought to have $F(f) = m$ for some m. However, the only path through $T_{\widehat{F}}$ compatible with f includes the node $?q$, and since $f(q)$ is undefined, the action of $T_{\widehat{F}}$ on f cannot yield m. Thus every question appearing within $T_{\widehat{F}}$ is within Q after all.

It is now easy to extract an enumerable graph for $F : R \to \mathbb{N}$ from $T_{\widehat{F}}$: each path $?q_0, p_0, \ldots, ?q_{r-1}, p_{r-1}, !m$ gives us a graph element $\langle\langle\langle q_0, p_0\rangle, \ldots, \langle q_{r-1}, p_{r-1}\rangle\rangle, m\rangle$. $\square$

In fact, the above argument can be adapted to work for many other sets $R \subseteq \mathsf{S}^1(Q)$. Let us say such an R is Δ-*semiseparable* if there is a Δ-enumerable set J_R of 1-codes such that for any 1-code a with $\mathrm{dom}\, a \subseteq Q$ we have

$$a \in J_R \text{ iff } R \cap U_a \text{ is inhabited.}$$

(In contrast to Definition 33, we do not here require a way of constructing an inhabitant of $R \cap U_a$ given a.)

Theorem 53 *If A is sequential, $Q \subseteq \mathbb{N}$, and $R \subseteq \mathsf{S}^1(Q)$ is Δ_A-semiseparable, then S^2_A is regular at Q, R.*

Proof As before, suppose $F \in \mathsf{S}^2(Q, R)$, take $\widehat{F} \in A(\overline{2})$ a realizer for F, and let $T_{\widehat{F}}$ be an irredundant Δ-enumerable sequential tree corresponding to $\widehat{F}$. If R is semiseparable, we may filter the nodes of $T_{\widehat{F}}$ to obtain a subtree $T'_{\widehat{F}}$ with the property that the path to any node is matched by some $f \in R$. We now claim that for every node $?q$ within $T'_{\widehat{F}}$ we have $q \in Q$. Otherwise, take some node $?q$ of minimal depth such that $q \notin Q$, take $f \in R$ matching the path to this node, and let $m = F(f)$. Then f is realized by some normalized $f' \in \Delta_\perp$ with $f' \upharpoonright_Q = f$. Define f'' by $f''(q) = \perp$ and $f''(j) = f'(j)$ elsewhere; then $f'' \in \Delta_\perp$ and f'' is also a realizer for f, so we ought to have $\widehat{F} \cdot f'' = \widehat{m}$. But this gives a contradiction, since the path through $T'_{\widehat{F}}$ determined by f'' goes through the node $?q$. Thus every question within $T'_{\widehat{F}}$ is in Q, and as before, it is now straightforward to extract a graph $g \in \Delta$ for F from $T'_{\widehat{F}}$. $\Box$

Even in sequential models, however, there are limits to the regularity phenomenon. We give here an artificial example of a pair Q, R at which regularity fails in *all* effective models. By the *subdomain* of a graph g, we shall mean the union of the domains of all 1-codes a such that some $\langle a, m \rangle$ appears in g.

Counterexample 54 Suppose A is any model in which $\Delta_A = \mathbb{N}^{\mathbb{N}}_{\mathrm{eff}}$. Let Q be the complement of the halting set including 0, and let

$$R = \{ f : Q \to \mathbb{N} \mid f \text{ computable}, \ f(0) \in Q \} ,$$

and define $F : R \to \mathbb{N}$ by $F(f) = f(f(0))$. Then F is clearly realized by $\lambda^* f.\, f(f\widehat{0})$. However, F cannot have a computable graph: if it did, its subdomain would be c.e., whereas it is clear that this must be the whole of Q which is not c.e..

Finally, we state without proof two results indicating that even when regularity fails at some $Q, \mathsf{S}^1(Q)$ (as in the case of $\mathsf{PC}^{\mathrm{eff}}$ and K_1), the 2-spectrum here nevertheless enjoys some interesting and perhaps surprising properties:

Theorem 55 *Suppose A is regular at $\mathbb{N}, \Delta$. Then for every $F \in \mathsf{S}^2(Q, \mathsf{S}^1(Q))$, there is an enumerable subset $Q' \subseteq Q$ such that if $f_0, f_1 \in \mathsf{S}^1(Q)$ and $f_0 \upharpoonright_{Q'} = f_1 \upharpoonright_{Q'}$, then $F(f_0) = F(f_1)$.*

Theorem 56 *Assume A satisfies Continuity, Enumeration, Restriction and P-Normalizability. Then for arbitrary $Q \subseteq \mathbb{N}$, any $F \in \mathsf{S}^2(Q, \mathsf{S}^1(Q))$ has a graph,* not necessarily in Δ_A, *with an enumerable subdomain $Q' \subseteq Q$.*

6 Non-deterministic models

We now present some mild generalizations of our main results, principally designed to cover models of *non-deterministic* computation, though they also have other incidental benefits.

As a first step, we shall relax our representation of natural numbers within a model to allow several elements of $A(\mathtt{N})$ to represent the same number n. Even for some deterministic settings this relaxation is useful: for example, in a time-sensitive programming language, there may be many ground type terms that all evaluate to the same natural number n but are nevertheless observationally distinct, and hence need to be represented by different elements of A.

Technically, in place of a system of numerals $\widehat{0}, \widehat{1}, \ldots$, we now postulate a modest set N with underlying set $\mathbb{N}$ and realizer type $\mathtt{N}$. We shall henceforth write $\Vdash$ for the realizability relation of N: thus, for every n there exists some $a \Vdash n$, and if $a \Vdash n$ and $a \Vdash n'$ then $n = n'$. In place of the conditions of Definition 2, we now require that there are modest set elements

$$suc \ \in \ (N \Rightarrow N)\ , \qquad rec \ \in \ (N \Rightarrow (N \Rightarrow N \Rightarrow N) \Rightarrow N \Rightarrow N)\ .$$

satisfying the relevant equations. Likewise, we shall now say our model A has ground-type iteration if there exists $iter : \overline{1} \to \overline{1}$ such that whenever $f \in A(\overline{1})$ and $a \Vdash n$ we have

$$\begin{array}{rcll} iter \cdot f \cdot a & \Vdash & m & \text{if } f \cdot a \ \Vdash \ inl(m) \\ iter \cdot f \cdot a & = & iter \cdot f \cdot b & \text{if } f \cdot a \ \Vdash \ inr(m) \text{ and } b = outr \cdot (f \cdot a)\ , \end{array}$$

where $outr$ is some fixed choice of realizer for the partial function $inr(m) \mapsto m$.

We continue to write Δ_A or just Δ for the set of functions $\mathbb{N} \to \mathbb{N}$ realizable in $A(\overline{1})$ (that is, the set of modest set morphisms $N \to \mathbb{N}$). We also write $e \Downarrow$ to mean that e evaluates to a realizer for some $n \in N$, and $e \Uparrow$ for $\neg\, e \Downarrow$.

The formulations of our axioms can now be readily adapted to this setting as follows:

For Continuity, we may work with the sets

$$\begin{array}{rcl} \bot & = & \{\alpha \mid a \Vdash j \Rightarrow \alpha \cdot a \Vdash 0\,\} \\ \top & = & \{\alpha \mid \exists t.\ (\forall j < t, a : \mathtt{N}.\ a \Vdash j \Rightarrow \alpha \cdot a \Vdash 0) \wedge (\forall a : \mathtt{N}.\ a \Vdash t \Rightarrow \alpha \cdot a \Vdash 1)\,\} \end{array}$$

and take as our Continuity Axiom the statement: For any $F \in A(\overline{2})$, if $F \cdot \alpha \Vdash n$ for all $\alpha \in \bot$, then $F \cdot \alpha \Vdash n$ for some $\alpha \in \top$.

For Enumeration, let us say that $f \in A(\overline{1})$ *strictly realizes* a partial function $h : \mathbb{N} \rightharpoonup \mathbb{N}$ if $a \Vdash n \in \mathrm{dom}\, h$ implies $f \cdot a \Vdash h(n)$, and $a \Vdash n \notin \mathrm{dom}\, h$ implies $f \cdot a \Uparrow$. We may now take our Enumeration Axiom to be the statement: For every $f \in A(\overline{1})$ strictly realizing some h, there exists $g \in \Delta_A$ whose offset range is the set of $\langle n, m\rangle$ such that $h(n) = m$.

T-Normalizability adapts very easily to the new setting: we again say that there is an element $norm : \overline{1} \to \overline{1}$ that transforms any realizer for any $f : N \to N$ into a canonical realizer for f. P-Normalizability can be likewise adapted using the notion of strict realizer, though we prefer to leave this axiom to one side for now.

For Collection, we now understand the *proper range* of $\Phi \in A(\sigma_0 \to \cdots \to \sigma_{r-1} \to \mathtt{N})$ to be the set of $n \in \mathbb{N}$ such that $\Phi \cdot x_0 \cdot \ldots \cdot x_{r-1} \Vdash n$ for some $x_0, \ldots, x_{r-1}$. Our axiom will now say that for any such Φ there exists $f \in A(\overline{1})$ strictly realizing some $h : \mathbb{N} \to \mathbb{N}$ whose range equals the proper range of Φ.

Finally, Restriction adapts very easily: we require that there is an element $\upharpoonright$ such that $a \upharpoonright b \Downarrow$ iff a, b are both realizers for natural numbers, and in this case $a \upharpoonright b$ realizes the same number as a.

It is now routine to verify that with these adaptations to the axioms, the proofs of Theorems 36, 38 and 42 go through in this more general setting with only bureaucratic changes (mostly just replacing each statement $e = \widehat{n}$ by $e \Vdash n$). This means that our results apply to time-sensitive models in which many observably distinct elements may evaluate to n, and also to Kleene's K_2 if we define our representation of natural numbers by $\beta \Vdash n$ iff $\beta(0) = n$: such a representation satisfies (the new version of) Continuity whereas a single-valued representation via numerals did not. This cures the blindspot regarding K_2 that we mentioned in Section 2.2.

We now come to the more interesting adaptation that allows for the possibility of non-determinism. We retain the modest set N and the realizability relation $\Vdash$ as above, but we also add a new relation $\Vdash^\diamond \subseteq A(\overline{1}) \times \mathbb{N}$. Informally, we read $a \Vdash n$ as 'a must yield the value n', and $a \Vdash^\diamond n$ as 'a may yield the value n'. For $\Vdash$ we retain the same conditions before as regards *suc*, *rec* and *iter*, but for $\Vdash^\diamond$ the only conditions we shall impose are:

1. $a \Vdash n$ implies $a \Vdash^\diamond n$,
2. $a \Vdash n$ and $a \Vdash^\diamond m$ imply $m = n$.

Of course, these conditions are satisfied if we take $\Vdash^\diamond = \Vdash$. In general, however, we may have $a \Vdash^\diamond n$ and $a \Vdash^\diamond m$ when $m \neq n$.

We shall continue to work mostly with $\Vdash$, which has all the same properties as before: indeed, the spectrum we wish to analyse will still consist of modest sets constructed from N by taking exponentials and regular subobjects. We therefore define the sets $\mathsf{S}^1(Q)$ and $\mathsf{S}^2(Q, R)$ as before, where $Q \subseteq \mathbb{N}$ and $R \subseteq \mathsf{S}^1(Q)$, and note that since R consists entirely of 'deterministic' functions, Proposition 29 still goes through as before: if $F : R \to \mathbb{N}$ has a graph in Δ the $F \in \mathsf{S}^2(Q, R)$.

The chief role played by $\Vdash^\diamond$ will be in a liberalized version of the Continuity Axiom that holds in typical non-deterministic models. Here we leave unchanged the above definitions of $\top, \bot$ in terms of the 'must' relation $\Vdash$, but we replace the usual axiom by:

Definition 57 (May-Continuity) *If A is a model with relations $\Vdash$ and $\Vdash^\diamond$ as above, then by the* May-Continuity Axiom *for A we shall mean the statement: For any $F \in A(\overline{2})$, if $F \cdot \alpha \Vdash^\diamond n$ for all $\alpha \in \bot$, then $F \cdot \alpha \Vdash^\diamond n$ for some $\alpha \in \top$.*

The intuitive justifications for Continuity mentioned in Section 2.2 also carry over to May-Continuity, in view of the fact that in typical cases we expect $\Vdash^\diamond$ to be a c.e. relation between 'programs' and natural numbers. Examples of models that satisfy May-Continuity but not ordinary Continuity include lattice models such as Scott's $\mathcal{P}\omega$ (where we take $a \Vdash n$ iff $a = \{n\}$, and $a \Vdash^\diamond n$ iff $n \in a$), and term models for programming languages with non-deterministic choice (where the 'must' and 'may' interpretations of $\Vdash$ and $\Vdash^\diamond$ are applied). Such models fell outside the scope of our investigation in [14], but we are hopeful that May-Continuity will now enable us to extend results of [14] to these models at all type levels.

We now see that May-Continuity alone is sufficient to support the classical KLS argument of Theorem 36. Assuming A to be a model equipped with relations $\Vdash, \Vdash^{\diamond}$ as above, we take $\top, \bot$ as above, and define $t : \top \cup \bot \to \mathbb{N}_\infty$ as usual. We also define the notion of a Δ-separable set as before, with reference only to $\Vdash$.

Theorem 58 *Assume A satisfies May-Continuity. Suppose $R \subseteq \Delta$ is Δ-separable and $F \in \mathsf{S}^2(\mathbb{N}, R)$. Then F is continuous on R: for any $f \in R$, there exists a 1-code a matching f such that $F(f') = F(f)$ for all $f' \in U_a \cap R$.*

Proof We explain briefly the necessary changes to the proof of Theorem 36. Again we take $\widehat{F}$ a realizer for F; notice that the hypothesis on F refers only to $\Vdash$, so that $\widehat{F}$ behaves deterministically on realizers $\widehat{f}$ for $f \in R$. We also use the notations $h, \zeta, \widehat{\zeta}, H$ as in the proof of Theorem 36, with the additional stipulation that each $\widehat{\zeta_i}$ is normalized (note that the functions ζ_i and their realizers $\widehat{\zeta_i}$ behave deterministically).

As before, for $\widehat{f}$ realizing $f \in \Delta$, $\alpha \in \top \cup \bot$ and $\widehat{m} \Vdash m \in N$, we define $\widehat{\xi}_{\widehat{f},\alpha,\widehat{m}}$ to behave as follows when applied to any $c \Vdash j \in \mathbb{N}$:

- If $j < t(\alpha)$, return a realizer for $f(j)$.
- Otherwise, search for $i \in H$ such that $\zeta_i(k) = f(k)$ for all $k < t(\alpha)$ and $F(\zeta_i) \neq m$.
- If such an i is found, return a realizer for $\zeta_i(j)$.

Note that $\widehat{\xi}_{\widehat{f},\alpha,\widehat{m}}$ is A-computable uniformly in $\widehat{f}, \alpha, \widehat{m}$.

We now say that $(\widehat{f}, \alpha, \widehat{m})$ *passes the may-test* if $\alpha \in \top$ and $\widehat{F} \cdot \widehat{\xi}_{\widehat{f},\alpha,\widehat{m}} \Vdash^{\diamond} m$. As before, $\langle a, m\rangle$ is a *suitable graph element* if $F(f) = m$ for all $f \in U_a \cap R$.

Claim 1: If $\widehat{f}$ realizes $f \in R$ and $(\widehat{f}, \alpha, \widehat{m})$ passes the may-test, then $\langle a, m\rangle$ is a suitable graph element, where $a = a(f, \alpha)$.

Proof: First, suppose for contradiction that $\zeta_i \in U_a \cap R$ and $F(\zeta_i) \Vdash m' \neq m$ (noting that $F(\zeta_i)$ has a deterministic value in $\mathbb{N}$). Then the search in the definition of $\widehat{\xi}_{\widehat{f},\alpha,\widehat{m}}$ must succeed in discovering some $\zeta_{i'}$ with $F(\zeta_{i'}) \Vdash m'' \neq m$, and then $\widehat{\xi}_{\widehat{f},\alpha,\widehat{m}}$ must represent $\zeta_{i'}$, so that $\widehat{F} \cdot \widehat{\xi}_{\widehat{f},\alpha,\widehat{m}} \Vdash m''$, contradicting $\widehat{F} \cdot \widehat{\xi} \Vdash^{\diamond} m$. So $F(\zeta_i) \Vdash m$ after all.

Next, consider a general $f' \in U_a \cap R$, and suppose $F(f') \Vdash m' \neq m$. Define f'_β as usual for $\beta \in \top \cup \bot$, and let $\widehat{f'}_\beta$ be a realizer for f'_β computed uniformly from β and a realizer for f'. Then for $\beta \in \bot$ we have $f'_\beta = f'$, whence $\widehat{F} \cdot \widehat{f'}_\beta \Vdash m'$ and so $\widehat{F} \cdot \widehat{f'}_\beta \Vdash^{\diamond} m'$. So by May-Continuity, we have $\widehat{F} \cdot \widehat{f'}_\beta \Vdash^{\diamond} m'$ for some $\beta \in \top$. But for each $\beta \in \top$ we have $f'_\beta \in R$ so that $F(f'_\beta)$ has a determinate value; hence $F(f'_\beta) = m'$, where $f'_\beta = \zeta_i$ for some i, contrary to what we showed above. This establishes Claim 1.

Claim 2: For any $\widehat{f}$ realizing $f \in R$ and any $\widehat{m} \Vdash m = F(f)$, there's some $\alpha \in \top$ such that $(\widehat{f}, \alpha, \widehat{m})$ passes the may-test.

Proof: For any $\alpha \in \bot$ we have that $\widehat{\xi}_{\widehat{f},\alpha,\widehat{m}}$ represents f where $F(f) = m$, so $\widehat{F} \cdot \widehat{\xi}_{\widehat{f},\alpha,\widehat{m}} \Vdash^{\diamond} m$. It follows immediately by May-Continuity that there is some $\alpha \in \top$ such that $\widehat{F} \cdot \widehat{\xi}_{\widehat{f},\alpha,\widehat{m}} \Vdash^{\diamond} m$, i.e. $(\widehat{f}, \alpha, \widehat{m})$ passes the may-test.

The theorem itself now follows immediately from Claims 1 and 2 as before. □

The argument of Section 4.2 now continues by invoking Collection to gather together the graph elements generated from all possible choices of $\widehat{f} \in A(\overline{1})$ and $\alpha \in A(\overline{1})$, using Restriction to filter out the unsuitable choices. In the non-deterministic setting there is a problem with this strategy: amongst all possible candidates for $\widehat{f}, \alpha \in A(\overline{1})$ will be some that behave non-deterministically, and in this case the behaviour of $\widehat{\xi}_{\widehat{f},\alpha,\widehat{m}}$ will be quite unpredictable. Moreover, there is no general way to filter out these non-deterministic values by means of Restriction tests. We are therefore unable to obtain the counterpart of Theorem 38 in the non-deterministic setting.

However, we are able to make further progress in the presence of T-Normalizability. Even here, a problem arises if we try to adapt the proof of Theorem 42 directly: there we made use of a realizer $\widehat{F}$ whose proper domain is some set $R^+ \subseteq R$, but there is now the possibility that $\widehat{F}$ may behave non-deterministically on $R^+ - R$, which will again lead to unpredictable behaviour for the $\widehat{\xi}_{\widehat{f},\alpha,\widehat{m}}$. We therefore restrict attention to separable domains R in order that our functions ζ_i may all be taken to lie within the original domain R of F.

A further point is that our construction only gives us computations that *may* yield suitable graph elements. We therefore need a version of Enumeration suitable for use with possibly non-deterministic functions:

Definition 59 (May-Enumeration) *By the* May-Enumeration Axiom *for A we shall mean the statement: For every $f \in A(\overline{1})$ there exists $g \in \Delta_A$ whose offset range is the set of $\langle n, m\rangle$ such that for some $a \Vdash n$ we have $f \cdot a \Vdash^\diamond m$.*

This stronger form of Enumeration is indeed valid for the kinds of non-deterministic models we have in mind, as the set of all possible values returned by a computation will typically be computably enumerable.

With these refinements in place, we may now obtain a regularity theorem in the T-Normalizable but non-deterministic setting. Here, as in Section 4.3, it is convenient to work with a cut-down version of the sets $\top, \bot$: that is, we consider $\bot_1 = \{\alpha_\infty\}$ and $\top_1 = \{\alpha_n \mid n \in \mathbb{N}\}$, where for each $n \in N_\infty$, $\alpha_n \in A(\overline{1})$ is the unique normalized element such that if $\widehat{i} \Vdash i$ then $\alpha_n \cdot \widehat{i} \Vdash 0$ if $i < n$ and $\alpha_n \cdot \widehat{i} \Vdash 0$ if $i \geq n$. Again as in Section 4.3, it is easy to see that the May-Continuity Axiom of Definition 57 implies the corresponding property for $\bot_1, \top_1$, and that this property suffices for the proof of Theorem 58.

It is also useful to note that in the presence of T-Normalizability we may obtain a normalizer $norm'$ for $\Vdash$-realizers of natural numbers:

$$norm' \;=\; \lambda a.\, norm\,(\lambda b.a)\, a\ .$$

When m is a natural number, we shall henceforth write $\widehat{m}$ for the unique normalized realizer of m.

Theorem 60 *Suppose A satisfies May-Continuity, T-Normalizability and May-Enumeration. Then A is regular at $\mathbb{N}, R$ for all Δ-separable $R \subseteq \Delta$.*

Proof Suppose $F \in \mathsf{S}^2(\mathbb{N}, R)$. Then the hypotheses of Theorem 58 are satisfied, and the proof of Theorem 58 goes through with $\top_1, \bot_1$ in place of $\top, \bot$. We now continue from where this proof left off, adapting the argument from the proof of Theorem 42.

Claim: For any $f \in R$ where $F(f) = m$, there exist $\zeta_i \in R$ and $\alpha \in \top_1$ such that $(\widehat{\zeta_i}, \alpha, \widehat{m})$ passes the may-test, where $f \in U_a$ for $a = a(\zeta_i, \alpha)$.

Proof: Given f, m as above and $\widehat{f}$ a normalized realizer for f, by Claim 2 in the proof of Theorem 58 we may take $\alpha \in \top_1$ so that $(\widehat{f}, \alpha, \widehat{m})$ passes the may-test and $f \in U_{a(f,\alpha)}$. Again as in the proof of Theorem 36, we may construct f_β for each $\beta \in \top_1 \cup \bot_1$ such that $f_\beta = f$ when $\beta \in \bot_1$, f_β is one of the ζ_i when $\beta \in \top_1$, and $f_\beta \in U_{a(f,\alpha)}$ for all β. Then $(\widehat{f_\beta}, \alpha, \widehat{m})$ passes the may-test for all $\beta \in \bot_1$, so by May-Continuity there exists $\beta \in \top_1$ such that $(\widehat{f_\beta}, \alpha, \widehat{m})$ passes the may-test. Taking i such that $\zeta_i = f_\beta$, we also have $\widehat{\zeta_i} = \widehat{f_\beta}$ since both are normalized; thus $(\widehat{\zeta_i}, \alpha, \widehat{m})$ passes the may-test. Furthermore, since $f_\beta \in U_{a(f,\alpha)}$, we have $a(f_\beta, \alpha) = a(f, \alpha)$, so $f \in U_{a(\zeta_i,\alpha)}$ as required. This proves the claim.

This shows that the graph elements $\langle a(\zeta_i, \alpha), m\rangle$ arising from triples $(\widehat{\zeta_i}, \alpha, \widehat{m})$ that pass the may-test are sufficient to constitute a set graph for F. We now need to collect these together to form an actual graph in Δ. Construct within A an enumeration $(i_0, \alpha_0, m_0), (i_1, \alpha_1, m_1), \ldots$ of all triples (i, α, m) with $\alpha \in \top_1 \cup \bot_1$ and $\zeta_i \in R$, and use this to construct an element $h \in A(\overline{1})$ such that for any $b_j \Vdash j$ we have

$$h \cdot b_j \ \Vdash^{\diamond} \ m' \quad \text{iff} \quad \widehat{F} \cdot \widehat{\xi}_{\zeta_i, \alpha_i, \widehat{m}_i} \text{ where } (i, \alpha, m) = (i_j, \alpha_j, m_j) \ .$$

Applying May-Enumeration to this h, we obtain $g \in \Delta$ whose offset range consists of the pairs $\langle j, m'\rangle$ satisfying the above condition. By easy programming within A, and again using the enumeration $j \mapsto (i_j, \alpha_j, m_j)$, we may transform this into $g' \in \Delta$ whose offset range is the set of $\langle a(\zeta_i, \alpha), m\rangle$ such that $(\widehat{\zeta_i}, \alpha_i, \widehat{m})$ passes the may-test, and hence into an enumerated graph $g'' \in \Delta$ for F. $\Box$

The above theorem applies to all non-deterministic models of interest satisfying T-Normalizability, including Scott's $\mathcal{P}\omega$, a range of non-deterministic game models, and term models for languages with non-deterministic choice. Of course, the T-Normalizability condition still excludes highly intensional models, so that (for instance) a term model for a language with both non-determinism and time-sensitivity would still not be covered by our treatment.

We do not know whether adopting P-Normalizability leads to any further results of interest in the non-deterministic setting.

7 The envelope of a spectrum

We have been working so far with the notion of a spectrum as a class of types generated from $\mathbb{N}$ via just two operations: taking (regular) subobjects, and forming function types $- \Rightarrow \mathbb{N}$. However, the intention behind the concept is that these may be taken to be representative of a wider class of types, built up from $\mathbb{N}$ via the formation of general function spaces, subobjects and quotients. In this section we justify this point of view by showing how a category of types with these closure properties may be reconstructed from an (abstract) spectrum, much as (under mild conditions) the whole of a simply-typed λ-algebra is recoverable from its pure type part (see [16, Section 4.2]). This will allow regularity results obtained for parts of the spectrum to be extended to a wider repertoire of types better suited to the needs of mathematical practice.

We have so far been dealing with spectra arising from TPCAs, but for our present purpose we need the notion of a spectrum in the abstract, of which the spectra of TPCAs furnish the leading examples. The formulation of the relevant closure conditions for an abstract spectrum requires a little effort:

Definition 61 *(i) An* (abstract) pre-spectrum S *is a tree whose nodes are labelled with sets (also known as* types*) such that:*

- *the root node is labelled with the set* $\mathsf{S}() = \mathbb{N}$*;*
- *if a node* X *is labelled with a set* $S = \mathsf{S}(Q_0, \ldots, Q_{k-1})$*, then for each subset* $Q_j \subseteq S$ *there is a branch from* X *to a node labelled with some set* $\mathsf{S}(Q_0, \ldots, Q_{k-1}, Q_k)$ *of functions from* Q_k *to* $\mathbb{N}$*, also written as* $Q_k \Rightarrow \mathbb{N}$*.*

The sets that label nodes are termed the proper types *of the spectrum, while the subsets of these (which label the edges) are called its* subtypes*. If* σ *is a subtype, we write* $\Pi(\sigma)$ *for the proper type from which* σ *arises as a subset.*

A proper type S *appearing in the spectrum as* $\mathsf{S}(Q_0, \ldots, Q_{k-1})$ *is said to be of* level k*, as are any subtypes* $R \subseteq S$*.*

(ii) The internal language *of a pre-spectrum* S *is defined by the following typing rules, where* σ, τ *range over subtypes of* S*, and* Γ *ranges over environments* $x_0 : \sigma_0, \ldots, x_{r-1} : \sigma_{r-1}$ *where the* x_i *are distinct variables:*

$$\frac{}{\Gamma \vdash x : \sigma}\; x : \sigma \in \Gamma \qquad \frac{\Gamma, x : \tau \vdash A : \mathtt{N}}{\Gamma \vdash \lambda x^{\tau}.A : \tau \to \mathbb{N}} \qquad \frac{\Gamma \vdash A : \tau \Rightarrow \mathbb{N} \quad \Gamma \vdash B : \tau}{\Gamma \vdash AB : \mathbb{N}}$$

$$\frac{\Gamma \vdash A : \sigma}{\Gamma \vdash |A| : \pi}\; \pi = \Pi(\sigma) \qquad \frac{\Gamma \vdash A : \pi}{\Gamma \vdash (A \upharpoonright \sigma) : \sigma}\; \pi = \Pi(\sigma)$$

$$\frac{\Gamma \vdash A : \mathbb{N}}{\Gamma \vdash \mathit{inl}\ A : \mathbb{N}} \qquad \frac{\Gamma \vdash A : \mathbb{N}}{\Gamma \vdash \mathit{inr}\ A : \mathbb{N}} \qquad \frac{\Gamma \vdash A : \mathbb{N} \quad \Gamma, z : \mathbb{N} \vdash B : \rho \quad \Gamma, z : \mathbb{N} \vdash C : \rho}{\Gamma \vdash (A\,?\,\mathit{inl}\ z \to B \mid \mathit{inr}\ z \to C) : \rho}$$

$$\frac{\Gamma \vdash A : \mathbb{N}}{\Gamma \vdash \mathit{fst}\ A : \mathbb{N}} \qquad \frac{\Gamma \vdash A : \mathbb{N}}{\Gamma \vdash \mathit{snd}\ A : \mathbb{N}} \qquad \frac{\Gamma \vdash A : \mathbb{N} \quad \Gamma \vdash B : \mathbb{N}}{\Gamma \vdash \langle A, B \rangle : \mathbb{N}} \qquad \frac{}{\Gamma \vdash n : \mathbb{N}}\; n \in \mathbb{N}$$

(iii) A valuation ν *for a type environment* Γ *is a function mapping each variable* $x : \sigma$ *in* Γ *to an element* $\nu(x)$ *of the set* σ*. If* $\Gamma \vdash A : \sigma$ *and* ν *is a valuation for* Γ*, we have a (possibly undefined)* interpretation $[\![A]\!]_\nu \in \sigma$ *defined by induction on typing derivations as follows:*

- $[\![x]\!]_\nu = \nu(x)$
- *If* $[\![A]\!]_{\nu, x \mapsto a}$ *is defined for all* $a \in \tau$*, and the function* $\Lambda a.\, [\![A]\!]_{\nu, x \mapsto a}$ *is present in* $\tau \to \mathbb{N}$*, then* $[\![\lambda x^{\tau}.A]\!]_\nu = \Lambda a.\, [\![A]\!]_{\nu, x \mapsto a}$*.*
- *If* $[\![A]\!]_\nu \in \tau \Rightarrow \mathbb{N}$ *and* $[\![B]\!]_\nu \in \tau$ *are both defined, then* $[\![AB]\!]_\nu = [\![A]\!]_\nu([\![B]\!]_\nu) \in \mathbb{N}$*.*
- *If* $[\![A]\!]_\nu \in \sigma$ *is defined and* $\pi = \Pi(\sigma)$*, then* $[\![|A|]\!]_\nu = [\![A]\!]_\nu \in \pi$*.*

- *If* $[\![A]\!]_\nu \in \pi$ *is defined,* $\pi = \Pi(\sigma)$ *and moreover* $[\![A]\!]_\nu \in \sigma$*, then* $[\![(A \upharpoonright \sigma)]\!]_\nu = [\![A]\!]_\nu$*.*
- *If* $[\![A]\!]_\nu \in \mathbb{N}$ *is defined then* $[\![inl\ A]\!]_\nu = inl([\![A]\!]_\nu)$ *where inl is some standard left injection* $\mathbb{N} \to \mathbb{N}$*. Likewise for inr, fst, snd and pairing.*
- *If* $[\![A]\!]_\nu = inl(n)$ *and* $[\![B]\!]_{\nu,z\mapsto n} = x \in \sigma$*, then* $[\![(A\,?\,inl\ z \to B \mid inr\ z \to C)]\!]_\nu = x$*; similarly for inr and C.*
- $[\![n]\!]_\nu = n$*.*

(iv) A pre-spectrum S *is an* (abstract) spectrum *if whenever* $\Gamma, x : \tau \vdash A : \mathtt{N}$*,* ν *is a valuation for* Γ *and* $[\![A]\!]_{\nu,x\mapsto a} \in \mathbb{N}$ *is defined for all* $a \in \tau$*, we have that* $[\![\lambda x^\tau.A]\!]_\nu$ *is defined (that is, the function* $\Lambda a.\,[\![A]\!]_{\nu,x\mapsto a}$ *is present in* $\tau \to \mathbb{N}$*).*

As regards the choice of injections $inl, inr : \mathbb{N} \to \mathbb{N}$, we shall for convenience assume $inl(0) = 0$ and $inr(0) = 1$. We shall also assume the chosen pairing operation $\mathbb{N} \times \mathbb{N} \to \mathbb{N}$ is bijective.

We now show how one may build a respectable category from an abstract spectrum. The idea is that the morphisms between types are precisely the functions definable in the internal language, possibly relative to other elements in the spectrum:

Definition 62 *If* σ, τ *are subtypes within a spectrum* S*, a function* $f : \sigma \to \tau$ *is called an* S-morphism *if there are a term* $\Gamma, x : \sigma \vdash A : \tau$ *and a valuation* ν *for* Γ *such that for every* $a \in \sigma$ *we have that* $[\![A]\!]_{\nu,x\mapsto a} \in \tau$ *is defined and equals* $f(a)$*.*

It is now just a matter of some lengthy calculations to show that this definition yields a category with good structure.

Proposition 63 *Let* S *be an abstract spectrum. The subtypes of* S *and* S*-morphisms between them form a concrete category, which we shall call the* pre-envelope $\mathcal{D}(\mathsf{S})$ *of* S*.*

Proof The identity function on any σ is defined by the term-in-context $x : \sigma \vdash x : \sigma$. If $f : \rho \to \sigma$ and $g : \sigma \to \tau$ are defined by terms $\Gamma, x : \rho \vdash A : \sigma$ and $\Delta, y : \sigma \vdash B : \tau$ in conjunction with valuations ν and μ respectively, where we may assume Γ, x and Δ, y are disjoint, then it is easy to check that $\Gamma, \Delta, x : \rho \vdash B[y \mapsto A] : \tau$ is a well-typed term which in conjunction with $\nu; \mu$ defines the composition $g \circ f$. The unit and associativity laws are inherited from those for ordinary sets and functions. □

Proposition 64 $\mathcal{D}(\mathsf{S})$ *has binary sums. They coincide with ordinary set-theoretic sums, and the level of a sum* $\sigma + \sigma'$ *is the maximum of the levels of* σ *and* σ'*.*

Proof We use the notation $A \succ \sigma$ to abbreviate the type coercion $(|A| \upharpoonright \sigma)$.

The following inductive steps suffice for constructing all sums $\sigma + \sigma'$. As part of the induction hypothesis, we include that fact that whenever $f : \sigma \to \rho$ and $g : \sigma' \to \rho$ are definable by terms $\Gamma, z : \sigma \vdash A[z] : \rho$ and $\Delta, z : \sigma \vdash B[z] : \rho$ in conjunction with valuations ν for Γ and μ for Δ, the resulting map $[f, g]_{\sigma,\sigma'} : \sigma + \sigma' \to \rho$ is definable by a term of the form

$$\Gamma, \Delta, x : \sigma + \sigma' \;\vdash\; (C[x]\ ?\ inl\, y \to A[D[x,y]] \mid inr\, y \to B[E[x,y]])\ : \rho$$

in conjunction with $\nu; \mu$, where

$$\begin{array}{rcl} x : \sigma + \sigma' & \vdash & C[x] : \mathbb{N} \\ \Gamma, x : \sigma + \sigma', y : \mathbb{N} & \vdash & D[x, y] : \sigma \\ \Delta, x : \sigma + \sigma', y : \mathbb{N} & \vdash & E[x, y] : \sigma' \end{array}$$

are all independent of f, g, A, B. We furthermore stipulate some semantic conditions on these contexts. Namely, if $b = inl_{\sigma,\sigma'}(a)$ then $[\![C[x]]\!]_{x \mapsto b} = inl(c)$ for some $c \in \mathbb{N}$, and $[\![D[x, y]]\!]_{x \mapsto b, y \mapsto c} = a$; likewise for inr and $E[x, y]$.

- The sum $\mathbb{N} + \mathbb{N}$ is $\mathbb{N}$, with injections $inl_{\mathbb{N},\mathbb{N}}$, $inr_{\mathbb{N},\mathbb{N}}$ defined by $x : \mathbb{N} \vdash inl\ x : \mathbb{N}$ and $x : \mathbb{N} \vdash inr\ x : \mathbb{N}$. For any pair of morphisms $f, g : \mathbb{N} \to \rho$ defined by $\Gamma, z : \mathbb{N} \vdash A : \rho$ and $\Delta, z : \mathbb{N} \vdash B : \rho$, the map $[f, g]_{\mathbb{N},\mathbb{N}} : \mathbb{N} + \mathbb{N} \to \rho$ is defined by $\Gamma, \Delta, x : \mathbb{N} \vdash (x\,?\,inl\, z \to A \,|\, inr\, z \to B) : \rho$. (Here and below, we suppress mention of the supporting valuations ν, μ.)

- If $\tau + \tau'$ has been defined where τ, τ' are proper types, then for any subtypes $\sigma \subseteq \tau$, $\sigma' \subseteq \tau'$ we may take $\sigma + \sigma'$ to be the evident subset of $\tau + \tau'$:

  $$\sigma + \sigma' \;=\; \{inl_{\tau,\tau'}(x) \mid x \in \sigma\} \;\cup\; \{inr_{\tau,\tau'}(x) \mid x \in \sigma'\}.$$

  If $inl_{\tau,\tau'}$ is definable by $\Gamma, x : \tau \vdash A[x] : \tau + \tau'$, then $inl_{\sigma,\sigma'}$ is definable by

  $$\Gamma, y : \sigma \vdash A[x \succ \tau] \succ \sigma + \sigma' \;\; : \sigma + \sigma\ ,$$

  and similarly for $inr_{\sigma,\sigma'}$. Given any pair of morphisms $f : \sigma \to \rho$, $g : \sigma' \to \rho$ defined by $\Gamma, z : \sigma \vdash A[z] : \rho$ and $\Delta, z : \sigma' \vdash B[z] : \rho$, take contexts $C[\,], D[\,], E[\,]$ appropriate for defining maps $[f', g']_{\tau,\tau'}$ as in the induction hypothesis; then the desired map $[f, g]_{\sigma,\sigma'} : \sigma + \sigma' \to \rho$ may be defined by

  $$\begin{array}{rll} \Gamma, \Delta, x : \sigma + \sigma' \;\vdash\; (C[x \succ \tau + \tau'] & ? & inl\, y \to A[D[x \succ \tau + \tau', y] \succ \sigma] \\ & | & inr\, y \to B[E[x \succ \tau + \tau', y] \succ \sigma']) \;: \rho\ , \end{array}$$

  Moreover, this term is itself clearly of the required form.

- Suppose σ is a subtype for which we have already constructed $\mathbb{N} + \sigma$, and $\tau = \sigma \Rightarrow \mathbb{N}$. Then we may take $\mathbb{N} + \tau$ to be a certain subtype of $(\mathbb{N} + \sigma) \Rightarrow \mathbb{N}$, namely the set of $e \in (\mathbb{N} + \sigma) \Rightarrow \mathbb{N}$ such that either

  - $e(inl_{\mathbb{N},\sigma}(inl_{\mathbb{N},\mathbb{N}}(n)) = 0$ for all $n \in \mathbb{N}$,
  - there exists $m \in \mathbb{N}$ such that $e(inl_{\mathbb{N},\sigma}(inr_{\mathbb{N},\mathbb{N}}(n)) = m$ for all $n \in \mathbb{N}$,
  - $e(inr_{\mathbb{N},\sigma}(x)) = 0$ for all $x \in \sigma$,

  or

  - $e(inl_{\mathbb{N},\sigma}(inl_{\mathbb{N},\mathbb{N}}(n)) = 1$ for all $n \in \mathbb{N}$,
  - $e(inl_{\mathbb{N},\sigma}(inr_{\mathbb{N},\mathbb{N}}(n)) = 0$ for all $n \in \mathbb{N}$.

The injection $inl_{\mathbb{N},\tau}$ takes $m \in \mathbb{N}$ to the function e of the first kind for this m, while $inr_{\mathbb{N},\tau}$ takes $d : \sigma \to \mathbb{N}$ to the function e of the second kind that behaves as d on its right component (we omit the tedious construction of the required formal terms). Furthermore, given morphisms $f : \mathbb{N} \to \rho$ and $g : \sigma \to \rho$ defined respectively by $\Gamma, z : \mathbb{N} \vdash A[z] : \rho$ and $\Delta, z : \sigma \vdash B[z] : \rho$, the desired map $[f,g]_{\mathbb{N},\tau}$ may clearly be defined by

$$\begin{array}{ll} \Gamma, \Delta, x : \mathbb{N}+\tau \ \vdash \ (\,|x|\,(inl_{\mathbb{N},\sigma}(0)) & ? \quad inl - \to A[\,|x|\,(inl_{\mathbb{N},\sigma}(1))] \\ & | \quad inr - \to B[\lambda y^{\sigma}.\,|x|\,(inr_{\mathbb{N},\sigma}(y))]\,) \ : \rho \end{array}$$

and this is of the required form.

Note that the three induction cases treated so far are enough to define $\mathbb{N}+\sigma$ for all σ.

- Finally, suppose that we have already constructed $\sigma + \sigma'$, and that $\tau = \sigma \Rightarrow \mathbb{N}$, $\tau' = \sigma' \Rightarrow \mathbb{N}$. We construct $\tau + \tau'$ as a subtype of $(\mathbb{N} + (\sigma + \sigma')) \Rightarrow \mathbb{N}$, namely the set of e of this type such that either

  - $e(inl_{\mathbb{N},\sigma+\sigma'}(n) = 0$ for all $n \in \mathbb{N}$,
  - $e(inr_{\mathbb{N},\sigma+\sigma'}(inr_{\sigma,\sigma'}(x))) = 0$ for all $x \in \sigma'$,

  or

  - $e(inl_{\mathbb{N},\sigma+\sigma'}(n) = 1$ for all $n \in \mathbb{N}$,
  - $e(inr_{\mathbb{N},\sigma+\sigma'}(inl_{\sigma,\sigma'}(x))) = 0$ for all $x \in \sigma$.

  The injection $inl_{\tau,\tau'}$ takes a function $d : \sigma \to \mathbb{N}$ to a function e of the first kind that behaves as d on the σ component; likewise for $inr_{\tau,\tau'}$ (we omit the formal details). Given $f : \sigma \to \rho$ and $g : \sigma' \to \rho$ defined respectively by $\Gamma, z : \sigma \vdash A[z] : \rho$ and $\Delta, z : \sigma' \vdash B[z] : \rho$, the desired map $[f,g]_{\tau,\tau'}$ may be defined by

$$\begin{array}{ll} \Gamma, \Delta, x : \mathbb{N}+\tau \ \vdash \ (\,|x|\,(inl_{\mathbb{N},\sigma}(0)) & ? \quad inl - \to A[\lambda y^{\sigma}.\,|x|\,(inr_{\mathbb{N},\sigma}(inl_{\sigma,\sigma'}(y)))] \\ & | \quad inr - \to B[\lambda y^{\sigma'}.\,|x|\,(inr_{\mathbb{N},\sigma}(inr_{\sigma,\sigma'}(y)))]\,) \ : \rho \end{array}$$

  which is of the required form.

By inspection of these constructions, we may also see by induction that our sums are isomorphic to ordinary disjoint unions, and that the level condition is satisfied. □

Proposition 65 $\mathcal{D}(\mathsf{S})$ *has finite products. They coincide with ordinary set-theoretic products, and the level of a product* $\sigma \times \sigma'$ *is the maximum of the levels of* σ *and* σ'.

Proof It is easy to check that the subtype $\{0\}$ of $\mathbb{N}$ serves as a terminal object.

The following inductive steps suffice for constructing all products $\sigma \times \sigma'$. We here omit syntactic details which are similar in flavour to those in the proof above.

- The product $\mathbb{N} \times \mathbb{N}$ is $\mathbb{N}$, with projections defined using the *fst* and *snd* term constructors, and the pairing of two morphisms $f, g : \rho \to \mathbb{N}$ constructed using $\langle -, - \rangle$.

- We take the product $\mathbb{N} \times (\sigma \Rightarrow \mathbb{N})$ to be a subtype of $(\mathbb{N} + \sigma) \to \mathbb{N}$, namely the set of $e \in (\mathbb{N} + \sigma) \to \mathbb{N}$ such that $e(inl_{\mathbb{N},\sigma}(n)) = e(inl_{\mathbb{N},\sigma}(n'))$ for all $n, n' \in \mathbb{N}$. The first projection gives the value of such an e on $inl_{\mathbb{N},\sigma}(0)$; the second projection gives the behaviour of e on the σ component. The pairing of morphisms $f : \rho \to \mathbb{N}$ and $g : \rho \to (\sigma \Rightarrow \mathbb{N})$ is an easy exercise.
- We may take the product $(\sigma \Rightarrow \mathbb{N}) \times (\sigma' \Rightarrow \mathbb{N})$ to be exactly the type $(\sigma + \sigma') \Rightarrow \mathbb{N}$, with the evident structure maps.
- Suppose τ, τ' are proper types and $\sigma \subseteq \tau$, $\sigma' \subseteq \tau'$ are subtypes of them. Having defined $\tau \times \tau'$, we may define $\sigma \times \sigma'$ to be the evident subset

  $$\{p \in \tau \times \tau' \mid fst_{\tau,\tau'}(p) \in \sigma,\ snd_{\tau,\tau'}(p) \in \sigma'\}\ .$$

  Projection and pairing for $\sigma \times \sigma'$ are now defined from those for $\tau \times \tau'$ with the help of type coercion, much as in the second induction case of the previous proof.

It is again clear by inspection that our products coincide with set-theoretic ones and that the level condition is satisfied. □

Proposition 66 *$\mathcal{D}(\mathsf{S})$ has exponentials $\sigma \Rightarrow \sigma'$. Elements of $\sigma \Rightarrow \sigma'$ are in canonical bijection with morphisms $\sigma \to \sigma'$; moreover, the exponentials $\sigma \Rightarrow \mathbb{N}$ coincide with the types previously denoted by $\sigma \Rightarrow \mathbb{N}$.*

Proof The following inductive steps suffice for constructing all exponentials:

- For any subtype σ of S, we take the exponential $\sigma \Rightarrow \mathbb{N}$ to be the type $\sigma \Rightarrow \mathbb{N}$ in the sense of Definition 61. There is then an evaluation morphism $(\sigma \Rightarrow \mathbb{N}) \times \sigma \to \mathbb{N}$ defined by $p : (\sigma \Rightarrow \mathbb{N}) \times \sigma \vdash (Fst[p])(Snd[p])$, where $Fst[p]$ and $Snd[p]$ are terms defining the projections for $(\sigma \Rightarrow \mathbb{N}) \times \mathbb{N}$. Given $f : \rho \times \sigma \to \mathtt{N}$ defined by $\Gamma, p : \rho \times \sigma \vdash A[p] : \mathtt{N}$, the transpose $\widetilde{f} : \rho \to (\sigma \Rightarrow \mathtt{N})$ is defined by $\Gamma, x : \rho \vdash \lambda y^{\sigma}.\, A[Pair_{\rho,\sigma}[x, y]]$, where $Pair_{\rho,\sigma}[\,]$ is a term denoting pairing of elements for $\rho \times \sigma$.
- For any σ and σ', we take $\sigma \Rightarrow (\sigma' \Rightarrow \mathbb{N})$ to be exactly the type $(\sigma \times \sigma') \Rightarrow \mathbb{N}$. The definitions of evaluation and transposition are straightforward in this case.
- Suppose $\sigma \Rightarrow \tau$ has been defined where τ is a proper type, and let ev denote its evaluation morphism. If $\sigma' \subseteq \tau$, let $\sigma \Rightarrow \sigma'$ be the subtype of $\sigma \Rightarrow \tau$ consisting of all e such that $ev(\langle e, x\rangle_{\sigma \Rightarrow \tau, \sigma}) \in \sigma'$ for all $x \in \sigma$. It is then easy to see how a term for ev yields one for evaluation at $\sigma \Rightarrow \sigma'$ with the help of some type coercion, and that the appropriate transpose construction goes through.

Finally, elements of $\sigma \Rightarrow \sigma'$ clearly correspond to morphisms $1 \to (\sigma \Rightarrow \sigma')$, since for any $f \in \sigma \Rightarrow \sigma'$ the term $z : \sigma \Rightarrow \sigma, x : 1 \vdash z : \sigma \Rightarrow \sigma'$ in conjunction with the valuation $z \mapsto f$ clearly defines $\Lambda * . f$. But the latter correspond to morphisms $\sigma \to \sigma'$ for general abstract reasons. □

Note that $\mathcal{D}(\mathsf{S})$ has all the subset types one could wish for, but not yet all quotient types. Indeed, this remark can be made precise via the following general definition.

Definition 67 *Suppose* $\mathbf{C}$ *is a category equipped with a functor* $I : \mathbf{C} \to \mathbf{Set}$.

(i) We say $(\mathbf{C}, I)$ has subobjects *if for every* $X \in \mathbf{C}$*, every subset inclusion* $\iota : S \to I(X)$ *can be lifted to a morphism* $\bar{\iota} : Y \to X$ *with* $I(\bar{\iota}) = \iota$*, such that for any morphism* $f : Z \to X$*, if* $I(f)$ *factors through* ι *then* f *factors through* ι'.

(ii) Dually, we say $(\mathbf{C}, I)$ has quotients *if for every* $X \in \mathbf{C}$*, every quotient map* $\delta : I(X) \to T$ *can be lifted to some* $\overline{\delta} : X \to Y$ *with* $I(\overline{\delta}) = \delta$*, such that for any* $f : X \to Z$*, if* $I(f)$ *factors through* δ *then* f *factors through* δ'.

We may now summarize the findings of the last few propositions as follows:

Theorem 68 *The category* $\mathcal{D}(\mathsf{S})$ *is cartesian closed and has binary sums, and the forgetful functor* $I : \mathcal{D}(\mathsf{S}) \to \mathbf{Set}$ *preserves finite products and sums. Furthermore,* $(\mathcal{D}(\mathsf{S}), I)$ *has subobjects.* □

The next step is to extend $\mathcal{D}(\mathsf{S})$ to a category with both subobjects and quotients. This is accomplished easily as follows:

Definition 69 *Given* $\mathcal{D}(\mathsf{S})$ *as above, define a category* $\mathcal{E}(\mathsf{S})$ *as follows:*

- *Objects are sets* $\sigma/\!\sim$*, where* σ *is an object of* $\mathcal{D}(\mathsf{S})$*,* $\sim$ *is a (total) equivalence relation on the set* σ*, and* $\sigma/\!\sim$ *is the set of equivalence classes.*
- *Morphisms* $\sigma/\!\sim\, \to \sigma'/\!\sim'$ *are ordinary functions* $\sigma/\!\sim\, \to \sigma'/\!\sim'$ *that are represented by some morphism* $\sigma \to \sigma'$ *in* $\mathcal{D}(\mathsf{S})$.

We write $J : \mathcal{E}(\mathsf{S}) \to \mathbf{Set}$ *for the evident forgetful functor.*

Theorem 70 $\mathcal{E}(\mathsf{S})$ *is a cartesian closed category which (with its forgetful functor* J*) has subobjects and quotients.*

Proof That $\mathcal{E}(\mathsf{S})$ is a category is straightforward. The terminal object is obvious, and the binary product $(\sigma/\sim) \times (\sigma'/\sim)$ is simply $(\sigma \times \sigma')/(\sim \times \sim')$. For the exponential $(\sigma/\sim) \Rightarrow (\sigma'/\sim')$, we let τ be the subtype of $\sigma \Rightarrow \sigma'$ consisting of all f such that $x \sim y$ implies $f(x) \sim' f(y)$, and define an equivalence relation $\approx$ on τ by: $f \approx g$ iff for all $x, y \in \sigma$, $x \sim y$ implies $f(x) \sim' g(y)$. It is routine to check that $(\tau/\approx)$ serves as the required exponential.

It is also easy to see that $\mathcal{E}(\mathsf{S})$ has subobjects and quotients: a subobject of a quotient of σ is just a quotient of a certain subobject of σ, and a quotient of a quotient of σ is again a quotient of σ. We leave the details to the reader. □

We shall refer to $\mathcal{E}(\mathsf{S})$ as the *envelope* of S: it unfolds what is implicit in S to yield a more generous realm of types suited to the demands of mathematical practice. It now remains to verify how this abstract story plays out in the concrete case of the spectrum derived from a computability model A.

Theorem 71 *Suppose* A *is any TPCA with weak numerals.*

(i) The spectrum S_A *of* A*, constructed as in Definition 28, is an abstract spectrum in the sense of Definition 61.*

(ii) In this case, $\mathcal{E}(\mathsf{S}_A)$ *is equivalent (over* $\mathbf{Set}$*) to the full subcategory of* $\mathcal{M}od(A)$ *generated from the modest set* N *of natural numbers via products, exponentials, subobjects and quotients.*

(Part (ii) of the theorem makes sense because $\mathcal{M}od(A)$ is also a concrete category possessing subobjects and quotients in the sense of Definition 67).

Proof (i) That S_A is a pre-spectrum is obvious. To verify the required closure property, we show that each typing judgement $\Gamma \vdash C : \tau$, where $\Gamma = x_0 : \sigma_0, \ldots, x_{r-1} : \sigma_{r-1}$, gives rise to an element $[C]_\Gamma \in A(\sigma_0 \to \cdots \to \sigma_{r-1} \to \tau)$, with the property that $[C]_\Gamma \cdot \nu(x_0) \cdot \ldots \cdot \nu(x_{r-1}) = [\![C]\!]_\nu \in \tau$ whenever the right-hand side is defined. Indeed, the definition of $[C]_\Gamma$ is straightforward by induction on the derivation of $\Gamma \vdash C : \tau$, using combinatory completeness for the case of λ-abstraction. Note also that we take $[|C|]_\Gamma = [C]_\Gamma$ and $[(C \upharpoonright \sigma)]_\Gamma = [C]_\Gamma$, so that type coercions are ignored by $[C]_\Gamma$. This suffices to show that if $[\![C]\!]_{\nu,x\mapsto a}$ is defined for all $a \in \tau$ then $\Lambda a.\, [\![C]\!]_{\nu,x\mapsto a}$ is present in the type $\tau \to \mathbb{N}$, since the latter by definition consists of all modest set morphisms $\tau \to \mathbb{N}$ that are realizable in A.

(ii) Every subtype σ within S_A actually arises as the underlying set of some modest set within R_A, so we may identify σ with this modest set. The translation $[C]_\Gamma$ above shows that every S_A-morphism $\sigma \to \sigma'$ is also a morphism of modest sets. Moreover, by inspecting in turn the above constructions for sums, products and exponentials within $\mathcal{D}(\mathsf{S}_A)$, one may verify that these correspond to sums, products and exponentials within $\mathcal{M}od(A)$, e.g. by showing that they are isomorphic to standard presentations of these constructions on modest sets. Since, in $\mathcal{M}od(A)$, elements of $X \Rightarrow Y$ correspond exactly to morphisms $X \to Y$, we may also read off from Proposition 66 that the modest set morphisms $\sigma \to \sigma'$ are *precisely* the S_A-morphisms. This amounts to showing that $\mathcal{D}(\mathsf{S}_A)$ is equivalent to a certain full subcategory of $\mathcal{M}od(A)$, all of whose objects may be generated from N by taking exponentials and subobjects.

Since also $\mathcal{M}od(A)$ has quotients, it follows easily that the whole of $\mathcal{E}(\mathsf{S}_A)$ is equivalent to a full subcategory $\mathcal{E}'(A)$ of $\mathcal{M}od(A)$, all of whose objects may be generated from N by exponentials, subobjects and quotients. Finally, since $\mathcal{E}(\mathsf{S}_A)$ is closed under products, exponentials, subobjects and quotients, and the constructions of these agree with those in $\mathcal{M}od(A)$, we conclude that $\mathcal{E}'(A)$ is precisely the full subcategory generated from N by these constructions. □

7.1 An application to computable analysis

Of course, our goal in establishing all of this is to be able to extend results about portions of S_A to other parts of the envelope. For example, if the spectra of two computability models A, B agree on some class $\mathbf{C}$ of types (for instance, on $\mathsf{S}(\mathbb{N}, R)$ for all effectively separable R), then $\mathcal{M}od(A), \mathcal{M}od(B)$ will also agree on all types that can be presented as subquotients of types in $\mathbf{C}$ in a uniform way. Rather than formulating a general theorem to this effect, we shall illustrate the idea with a typical example.

Let $\mathbb{D}$ be a subset of $\mathbb{R}^n$ for some n. We ask to what extent the set of A-computable functions $\mathbb{D} \to \mathbb{R}$ is robust with respect to the choice of A. To make this precise, we may consider the (Cauchy) real number object R within $\mathcal{M}od(A)$, and the regular subobject $D \subseteq R^n$ determined by $\mathbb{D} \cup |R^n|$; we are interested in the modest set morphisms $D \to R$.

It is easy to see that R is isomorphic within $\mathcal{M}od(A)$ to a quotient of $\mathbb{N}^\mathbb{N}$ by an equivalence relation that can be defined purely set-theoretically, without reference to A. It follows easily, either by our general results above or by a simple bespoke construction, that R^n is also isomorphic to a certain quotient of $\mathbb{N}^\mathbb{N}$, and hence that D is isomorphic to a quotient of a subobject of $\mathbb{N}^{\mathbb{N}}$, defined with reference to $\mathbb{D}$ but without reference to A. Thus, D in effect appears as a subtype of level 1 within the spectrum of A, and $D \Rightarrow N$ as a proper type of level 2.

It follows that if A satisfies the hypotheses of Section 4.2 (Continuity, Iterated Restriction, Collection, Enumeration) and $\mathbb{D} \subseteq \mathbb{R}^n$ is Δ_A-*separable* (in an evident sense coherent with Definition 33), then the morphisms $D \to N$ are precisely those that have 'graphs' within Δ_A; the set of such morphisms is thus determined solely by Δ_A. (As noted in Section 4.1, the class of separable sets is quite extensive and includes all enumerable unions of basic open or closed sets, along with less benign examples such as the set of rational points or its complement.) Likewise, if A satisfies the hypotheses of Section 4.3 (Continuity, T-Normalizability, Enumeration), the same is true without any condition on $\mathbb{D}$. Thus, for example, even for arbitrary $\mathbb{D}$, we obtain (for example) a fairly robust class of effectively computable functions $\mathbb{D} \cap \mathbb{R}^n_{\mathrm{eff}} \to \mathbb{R}_{\mathrm{eff}}$, where $\mathbb{R}_{\mathrm{eff}}$ is the usual set of computable reals.

This relatively simple example serves to illustrate the envisioned use of $\mathcal{E}(\mathsf{S}_A)$ as a realm of types that supporting typical constructions from ordinary mathematics. We expect more far-reaching examples of this kind to be forthcoming once we have extended our regularity theorems to level 3 and above.

8 Prospects for extensions and further work

8.1 Extensions to higher types

We now comment briefly on what will be involved in extending our results to type 3 and above, as we intend to do in Part II.

First, let us consider the task of recovering the results of [14] in our setting — that is, of analysing the contents of the *pure* types, with no subset types involved. As in Section 4, our task at each type level will be to show that the class of functionals with a *graph* in Δ provides both a lower and an upper bound for the set of functionals present. At type 2, the lower bound was relatively trivial, and the main work was to use the KLS machinery to establish the upper bound. At type 3, by contrast, the upper bound appears to flow rather easily from the work we have done already (much as the proof of the higher-type version of KLS follows relatively cheaply from the type 2 version), and it is the lower bound that causes the difficulty. That is to say, programming an operation that 'applies' a type 3 graph to a given type 2 functional is non-trivial, and it is here that the ingenious construction of Normann [18] comes in. Nevertheless, we fully expect that at least in the T-Normalizable setting, a relatively straightforward adaptation of this construction (and its elaboration in [14]) will suffice to complete the analysis at all pure types: the result will be that the total functionals in A are the 'Kleene-Kreisel functionals relativized to Δ_A' (cf. [16, Section 9.5]). Note, however, that at this point the setting of a TPCA with iteration is no longer adequate: we need to assume the presence of a *recursion* operator (as we did in [14]) in order to implement the Normann programs in A. (That this is essential is shown by the left-bounded sequential procedure model, which does not yield the full class of continuous functionals at type 3.)

The situation is less clear in the non-normalizable case. The hope is that we can adapt the arguments from [14] for the effective case, which were designed to get round the absence of normalizability. However these arguments were in some places quite subtle, and their translation to the axiomatic setting is not completely evident. A particular challenge will be to find a counterpart for Proposition 8.10 of [14] which involved a delicate diagonal construction. It is this challenge that our Dependent Collection Axiom (Definition 10) is designed to meet, although whether this suffices remains to be seen.

There is also the question of how far into the rest of the spectrum our results can be extended, once subset types are allowed. The situation here potentially becomes increasingly complex as we pass to higher types: for instance, in order to characterize some type $\mathsf{S}^3(Q, R, S)$, it may be that conditions on any or all of Q, R, S are necessary. Although our expectation is that we will be able to account for enough of the spectrum to cover many types of interest for computable analysis, exactly how much to expect here is not yet clear.

8.2 Other questions

There are of course many further questions clamouring for attention. We mention here a selection:

- To complete the reconstruction of results of [14] in our setting, one should also consider *relative realizability* interpretations of types (which are presumably straightforward) as well as *modified realizability* (which could present a more substantial challenge).
- Our analysis of subtypes has so far been confined to *regular* subobjects: those subobjects $Y \subseteq X$ such that the realizability structure on Y is simply the restriction to Y of the one on X. Can anything of interest be said for more general subobjects? A typical example for consideration here is the subtype of $(N \Rightarrow 2) \Rightarrow 2$ (interpreted over K_1) consisting of the *uniformly continuous* functionals: the *fan functional* will be computable on this subtype if the latter is taken to be a suitably defined non-regular subobject, but not if it is crudely interpreted as the corresponding regular subobject.
- Our work has suggested an extension of the Kleene-Kreisel functionals with subtypes and quotients, but we have so far completely neglected the topological side to the story. For example, can our extended category be naturally described in terms of limit spaces or other quasi-topological structure? It would also seem natural to revisit the work of Escardó [6] on questions of searchability and exhaustibility for predicates on subsets of simple types σ. Does the picture change if such predicates no longer have to be presented by total functions on the whole of σ?
- Once the appropriate higher-type extensions of our results have been worked out, we would like to look more deeply at what exactly this implies for computability on various spaces arising in analysis. For instance, we would expect the implications for examples 3 and 4 from Section 1.2 to be fairly clear: there should be robust computability notions at the relevant types, and indeed the operations in question should admit implementations in PCF that behave correctly even on inputs implemented in much richer languages. But we are also eager to explore what other phenomena arise in more advanced areas of analysis: the extensive body of existing work on computable analysis (e.g. [19, 21]) is likely to be relevant here. The relationship to existing work on higher types over the reals should also be clarified (see [20] and the references therein).

- Finally, we would also like to know what other portions of higher-type computability theory can be developed on the basis of the axioms we have given, or others like them. A particular candidate for attention is the general theory of *models of PCF* and their equational theories — the rudiments of an axiomatic treatment (drawing on experience from synthetic domain theory) were presented in [16, Section 7.1], but we believe there is scope for a more systematic development.

  We would regard the distillation of a simple set of axioms on TPCAs with rich and far-reaching consequences spanning several aspects of the theory as a very positive development for the subject as a whole.

References

[1] Abramsky, S., McCusker, G.: *Game semantics.* In: Schwichtenberg, H., Berger, U. (eds.), *Computational Logic: Proceedings of the 1997 Marktoberdorf Summer School*, pp. 1–56, Springer, Heidelberg (1999)

[2] Bauer, A.: *First steps in synthetic computability theory.* Electronic Notes in Theoretical Computer Science **155**, 5–31 (2006)

[3] Bucciarelli, A., Ehrhard, T., Manzonetto, G.: *Not enough points is enough.* In: Duparc, J., Henzinger, T.A. (eds.), *CSL07: Proceedings of 16th Computer Science Logic*, pp. 298–312, Springer (2007)

[4] Church, A.: *A formulation of the simple theory of types.* Journal of Symbolic Logic **5(2)**, 56–68 (1940)

[5] Cockett, R., Hofstra, P.: *Introduction to Turing categories.* Annals of Pure and Applied Logic **156(2-3)**, 183–209 (2008)

[6] Escardó, M.H.: *Exhaustible sets in higher-type computation.* Logical Methods in Computer Science **4(3)**, 37 pages (2008)

[7] Feferman, S.: *Theories of finite type related to mathematical practice.* In: Barwise, J. (ed.), *Handbook of Mathematical Logic*, pp. 913–941, Elsevier, Amsterdam (1977)

[8] Fenstad, J.E.: *General Recursion Theory.* Springer, Berlin (1980)

[9] Friedberg, R.M.: *Un contre-exemple relatif aux fonctionelles récursives.* Comptes Rendus de l'Académie des Sciences, Paris **247**, 852–854 (1958)

[10] *The HOL system logic*, Version 11 (2017). Available from https://hol-theorem-prover.org/

[11] Hyland, J.M.E.: First steps in synthetic domain theory. In: Carboni, A., Pedicchio, M.C., Rosolini, G. (eds.), *Category Theory: Proceedings of the International Conference in Como, 1990*, pp. 131–156. Springer, Berlin (1991)

[12] Kreisel, G.: *Some reasons for generalizing recursion theory*. In: Gandy, R.O., Yates, C.M.E. (eds.), *Logic Colloquium 69: Proceedings of the Summer School and Colloquium in Mathematical Logic, Manchester*, pp. 139–198. North-Holland, Amsterdam (1971)

[13] Kreisel G., Lacombe, D., Shoenfield, J.R.: *Partial recursive functionals and effective operations*. In: Heyting, A. (ed.), *Constructivity in Mathematics: Proceedings of the Colloquium held in Amsterdam, 1957*, pp. 290–297. North-Holland, Amsterdam (1959)

[14] Longley, J.: *On the ubiquity of certain total type structures*. Mathematical Structures in Computer Science **17(5)**, 841–953 (2007)

[15] Longley, J.: *Computability structures, simulations and realizability*. Mathematical Structures in Computer Science **24(2)**, 49 pages (2014)

[16] Longley, J., Normann, D.: *Higher-Order Computability*. Theory and Applications of Computability, Springer (2015)

[17] Moschovakis, Y.N.: *Axioms for computation theories — first draft*. In: Gandy, R.O., Yates, C.M.E. (eds.), *Logic Colloquium 69: Proceedings of the Summer School and Colloquium in Mathematical Logic, Manchester*, pp. 199–255. North-Holland, Amsterdam (1971)

[18] Normann, D.: *Computability over the partial continuous functionals*. Journal of Symbolic Logic **65(3)**, 1133–1142 (2000)

[19] Pour-El, M.B., Richards, I.: *Computability in Analysis and Physics*. Springer, Berlin (1989)

[20] Schröder, M.: $\mathbb{N}^{\mathbb{N}^{\mathbb{N}}}$ *does not satisfy Normann's condition*. ACM Computing Research Repository abs/1010.2396 (2010)

[21] Weihrauch, K.: *Computable Analysis: An Introduction*. Springer, Berlin (2000)