EVERY COUNTABLE MODEL OF SET THEORY EMBEDS INTO ITS OWN CONSTRUCTIBLE UNIVERSE

JOEL DAVID HAMKINS

Abstract. The main theorem of this article is that every countable model of set theory $\langle M, \in^M\rangle$, including every well-founded model, is isomorphic to a submodel of its own constructible universe $\langle L^M, \in^M\rangle$. In other words, there is an embedding $j : M \to L^M$ that is elementary for quantifier-free assertions. It follows from the proof that the countable models of set theory are linearly pre-ordered by embeddability: for any two countable models of set theory $\langle M, \in^M\rangle$ and $\langle N, \in^N\rangle$, either M is isomorphic to a submodel of N or conversely. Indeed, they are pre-well-ordered by embeddability in order-type exactly ω_1+1. Specifically, the countable well-founded models are ordered under embeddability exactly in accordance with the heights of their ordinals; every shorter model embeds into every taller model; every model of set theory M is universal for all countable well-founded binary relations of rank at most Ord^M; and every ill-founded model of set theory is universal for all countable acyclic binary relations. Finally, strengthening a classical theorem of Ressayre, the same proof method shows that if M is any nonstandard model of PA, then every countable model of set theory—in particular, every model of ZFC plus large cardinals—is isomorphic to a submodel of the hereditarily finite sets $\langle \mathrm{HF}^M, \in^M\rangle$ of M. Indeed, $\langle \mathrm{HF}^M, \in^M\rangle$ is universal for all countable acyclic binary relations.

1. Introduction

In this article, I shall prove that every countable model of set theory $\langle M, \in^M\rangle$, including every well-founded model, is isomorphic to a submodel of its own constructible universe $\langle L^M, \in^M\rangle$. Another way to say this is that there is an embedding

$$j : \langle M, \in^M\rangle \to \langle L^M, \in^M\rangle$$

that is elementary for quantifier-free assertions in the language of set theory.

Main Theorem 1. *Every countable model of set theory $\langle M, \in^M\rangle$ is isomorphic to a submodel of its own constructible universe $\langle L^M, \in^M\rangle$.*

The proof uses universal digraph combinatorics, including an acyclic version of the countable random digraph, which I call the countable random $\mathbb{Q}$-graded digraph, and higher analogues arising as uncountable Fraïssé limits, leading eventually to what I call the hypnagogic digraph, a set-homogeneous, class-universal, surreal-numbers-graded acyclic class digraph,

My research has been supported in part by NSF grant DMS-0800762, PSC-CUNY grant 64732-00-42 and Simons Foundation grant 209252. In addition, I gratefully acknowledge the generous support provided to me as a visiting fellow in Spring 2012 at the Isaac Newton Institute for Mathematical Sciences in Cambridge, U.K., where most of this research was undertaken. I thank Philip Welch for helpful discussions undertaken whilst strolling the Churchill College grounds in Cambridge and a math-filled excursion to Wells in southern England, with cream tea, a Saxon church and mathematics in Bradford-on-Avon. Thanks also to Ali Enayat, Victoria Gitman and Emil Jeřábek for helpful comments.

which is closely connected with the surreal numbers. The proof shows that $\langle L^M, \in^M \rangle$ contains a submodel that is a universal acyclic digraph of rank Ord^M, and so in fact this model is universal for all countable acyclic binary relations of this rank. When M is ill-founded, this includes all acyclic binary relations. The method of proof also establishes the following, thereby answering a question posed by Ewan Delanoy [Del11].

Main Theorem 2. *The countable models of set theory are linearly pre-ordered by embeddability: for any two countable models of set theory $\langle M, \in^M \rangle$ and $\langle N, \in^N \rangle$, either M is isomorphic to a submodel of N or conversely. Indeed, the countable models of set theory are pre-well-ordered by embeddability in order type exactly $\omega_1 + 1$.*

The proof shows that the embeddability relation on the models of set theory conforms with their ordinal heights, in that any two models with the same ordinals are bi-embeddable; any shorter model embeds into any taller model; and the ill-founded models are all bi-embeddable and universal.

The proof method arises most easily in *finite* set theory, showing that the nonstandard hereditarily finite sets HF^M coded in any nonstandard model M of PA or even of $I\Delta_0$ are similarly universal for all acyclic binary relations. This strengthens a classical theorem of Ressayre, while simplifying the proof, replacing a partial saturation and resplendency argument with a soft appeal to graph universality.

Main Theorem 3. *If M is any nonstandard model of* PA*, then every countable model of set theory is isomorphic to a submodel of the hereditarily finite sets $\langle \mathrm{HF}^M, \in^M \rangle$ of M. Indeed, $\langle \mathrm{HF}^M, \in^M \rangle$ is universal for all countable acyclic binary relations.*

In particular, every countable model of ZFC and even of ZFC plus large cardinals arises as a submodel of $\langle \mathrm{HF}^M, \in^M \rangle$. Thus, inside any nonstandard model of finite set theory, we may cast out some of the finite sets and thereby arrive at a copy of any desired model of infinite set theory, having infinite sets, uncountable sets or even large cardinals of whatever type we like.

A structure M is *universal* for a class Δ of structures, if every structure in Δ is isomorphic to a substructure of M. In this article I use the term *model of set theory* to mean a first-order structure $\langle M, \in^M \rangle$ satisfying at least the Kripke-Platek KP axioms of set theory, a very weak fragment of ZF. Although typically for the models of set theory I have in mind the models of ZFC or even ZFC plus large cardinals, it turns out that the results of this article go through for the much weaker theory KP, and indeed still weaker theories will suffice, but for definiteness I shall officially use KP, leaving the determination of how weak we may go for a later project.

The language of set theory has only the set-membership relation $\in$, and so a *submodel* of a model $\langle M, \in^M \rangle$ of set theory is simply a subset $N \subseteq M$, where one restricts the relation to form $\langle N, \in^M \restriction N \rangle$. An *embedding* of one model $\langle M, \in^M \rangle$ into another $\langle N, \in^N \rangle$ is an isomorphism of M with a submodel of N, that is, a function $j : M \to N$ for which $x \in^M y$ if and only if $j(x) \in^N j(y)$. Since a model of set theory is a set with a binary relation, it is technically a certain special kind of directed graph, and many of the arguments of this article will proceed from this graph-theoretic perspective. From this perspective, a submodel of a model of set theory, viewed as a directed graph, is the same as an induced subgraph.

The three main theorems reappear in this article as theorems 29, 33 and 9, respectively.

2. The countable random $\mathbb{Q}$-graded digraph

The main theorems will be proved by finding copies of certain universal digraphs among the submodels of the models of set theory under consideration. So let me begin by developing a little of this universal digraph theory. A *digraph*, or directed graph, is a structure $\langle G, \rightharpoonup\rangle$ where G is a set of vertices, or nodes, and $\rightharpoonup$ is a binary relation on G. A digraph is *acyclic* if there is no finite directed path from a vertex to itself. That is, an acyclic digraph is one with no directed cycles. Note that the undirected version of an acyclic digraph may be far from a tree, since there can be undirected cycles.

Let $\mathbb{Q}$ be the endless dense linear order of the rational numbers. A digraph G is $\mathbb{Q}$-*graded* if every vertex a in G is assigned a rational value q_a in such a way that $q_a < q_b$ whenever $a \rightharpoonup b$. Such a graph must be acyclic, since the values increase along any directed path. More generally, for any linear order ℓ we may consider the ℓ-graded directed graphs, and these also are acyclic.

Lemma 4. *Every countable acyclic digraph can be* $\mathbb{Q}$*-graded by a suitable value assignment.*

Proof. Suppose that G is a countable acyclic digraph. Let $\lhd$ be the reachability relation on G, the transitive closure of the edge relation of G. Because G is acyclic, it follows that $\lhd$ is a partial order on the vertices. Since every partial order extends to a linear order, and every countable linear order embeds into $\mathbb{Q}$, it follows that we may assign the vertices of G to rational numbers in such a way that $a \lhd b$ implies $q_a < q_b$. In particular, this is a $\mathbb{Q}$-grading of G. ☐

Although the value assignment arising in the proof of lemma 4 is one-to-one, in general a $\mathbb{Q}$-graded digraph may have many nodes with the same value, and the universal graded digraphs appearing later in this article will have infinitely many nodes of each value. The grading of a digraph is in effect a laying-out of its vertices on levels, in such a way that every directed edge points from a lower-level node to a higher-level node.

An alternative grading concept would regard a graded digraph $\langle G, \rightharpoonup, \leq\rangle$ as a digraph $\langle G, \rightharpoonup\rangle$ together with a linear pre-order $\leq$ on the vertices whose corresponding strict order includes the edge relation, so that $a \rightharpoonup b$ implies $a < b$. This formalization founds the graded digraph concept in a finite first-order language, but in effect allows embeddings and isomorphisms to act also on the linear order, preserving merely the relative value rather than the exact value; but I shall not need this, and so in this article I have opted for the value-assignment concept of grading above and insist that graded digraph embeddings preserve these values. A structure A is *homogeneous* if every isomorphism of finitely generated substructures of A extends to an automorphism of A.

Theorem 5. *There is a countable homogeneous* $\mathbb{Q}$*-graded digraph* Γ*, which is universal for all countable* $\mathbb{Q}$*-graded digraphs. Furthermore,*

(1) *There is a unique such digraph up to isomorphism, even when restricting universality to finite* $\mathbb{Q}$*-graded digraphs.*
(2) *Every countable acyclic digraph arises as an isomorphic copy of an induced subgraph of* Γ.
(3) Γ *admits a computable presentation.*

We call this Γ the *countable random* $\mathbb{Q}$*-graded digraph.*

Proof. There are several independent constructions of this highly canonical object. Let me begin with the abstract realization of Γ as a Fraïssé limit. Let $\mathcal{Q}$ be the collection of finite $\mathbb{Q}$-graded digraphs. Note that $\mathcal{Q}$ enjoys the *hereditary* property, that any induced subgraph of a graph in $\mathcal{Q}$ is in $\mathcal{Q}$, and also the *joint embedding* property, asserting that any two members of $\mathcal{Q}$ embed into a third. Furthermore, $\mathcal{Q}$ has the *amalgamation* property, meaning that whenever A, B and C are in $\mathcal{Q}$, where A embeds into B and also into C, then there is a D into which both B and C embed while agreeing on the image of A. These three properties are sufficient that $\mathcal{Q}$ has a Fraïssé limit Γ, a countable, homogeneous, directed $\mathbb{Q}$-graded graph, whose finite induced subgraphs are up to isomorphism precisely the graphs in $\mathcal{Q}$ (see [Hod93, thm 7.1.2]). The universal property of the Fraïssé limit ensures that Γ is universal for all countable acyclic digraphs, because any such graph G is $\mathbb{Q}$-graded by lemma 4 and since G is the direct limit of its finite induced subgraphs, each of which is in $\mathcal{Q}$, it follows that one may amalgamate the limit construction of G as occurring inside Γ.

Consider alternatively the following construction of Γ via forcing, which amounts essentially to the forcing construction of the Fraïssé limit. Let $\mathbb{P}$ be the partial order of all finite $\mathbb{Q}$-graded digraphs with nodes coming from a fixed countably infinite set, ordering the digraphs of $\mathbb{P}$ by the induced subgraph relation. If $A \subseteq B$ and C are any such finite $\mathbb{Q}$-graded digraphs, with an embedding $f : A \to C$, then let $D_{A,B,C,f}$ be the set of graphs $H \in \mathbb{P}$ such that if $C \subseteq H$, then f extends to an embedding $g : B \to H$. This is a dense collection of conditions H, since for any $H_0 \in \mathbb{P}$, if $C \subseteq H_0$ then we can extend H to include a copy of B over $f(A)$, and thus extend f to $g : B \to H$, as desired. Since there are only countably many such dense sets, we may build a chain of graphs $H_0 \subseteq H_1 \subseteq \cdots$ such that $\Gamma = \bigcup_n H_n$ meets every such dense set. It follows that Γ is *weakly homogeneous*, which means that whenever $A \subseteq B$ and $f : A \to \Gamma$, then there is an extension to $g : B \to \Gamma$. This implies that Γ is universal for all countable $\mathbb{Q}$-graded digraphs, since any such graph G is the union $\bigcup_n G_n$ of finite graphs, and weak homogeneity allows us to build up a chain of embeddings $f_n : G_n \to \Gamma$, which realizes G as an induced subgraph of Γ. It also implies that Γ is fully homogeneous, by a similar back-and-forth argument that systematically extends a finite partial isomorphism to a full automorphism of Γ.

Let me now describe a more concrete construction, which will realize Γ as a computable directed $\mathbb{Q}$-graded graph. We build Γ as the union of a sequence of finite directed $\mathbb{Q}$-graded graphs $\Gamma_n \in Q$. Namely, let Γ_0 be the empty graph, and at each stage n, we consider all possible ways to extend Γ_n by the addition of one vertex with a value amongst the first n rational numbers, and let Γ_{n+1} realize all such possibilities. Since we can carry out this process by a computational procedure, the limit $\Gamma = \bigcup_{n\in\omega} \Gamma_n$ is a computable directed $\mathbb{Q}$-graded graph. This graph Γ is universal for all countable directed cycle-free graphs G, since any such G can be $\mathbb{Q}$-graded, and then we can build up the embedding of G into Γ in stages. Namely, we enumerate the vertices of G as v_0, v_1, and so on, and at stage n we have a copy of the first n vertices of G inside some Γ_{k_n}. The next vertex v_n has a certain value and exhibits a certain pattern of connectivity to the prior vertices of G. At a stage in the construction of Γ where the value q_n of v_n is allowed, we add a vertex exhibiting precisely that pattern of connectivity over the copy of the prior vertices that we have built in Γ. Thus, we may extend our embedding to include a copy of v_n with the correct value. By continuing this process, we end with an embedding of G into Γ, as desired.

Finally, let me give a probabilistic account of Γ. Let Γ have infinitely many vertices of every rational value, and for each pair of nodes a and b with a having lower value than b, flip a coin to determine whether we place an edge between them. Thus, Γ is the random $\mathbb{Q}$-graded directed graph. With probability one, every finite $\mathbb{Q}$-graded directed graph arises as an induced subgraph, and furthermore, any particular finite subgraph A will find with probability one an extension inside Γ realizing any particular pattern of connectivity for a new vertex having a particular value (subject to the constraints on values). Thus, almost surely the resulting graph Γ will exhibit the finite-pattern property that allows us to extend any partial isomorphism via a back-and-forth argument to a full automorphism. □

The need to consider graded digraphs, rather than mere digraphs, arises from the fact that there simply are no nontrivial homogeneous acyclic digraphs:

Observation 6. *If an acyclic digraph G has $a \rightharpoonup b \rightharpoonup c$ as an induced subgraph, then it is not homogeneous as a digraph (without any grading structure).*

Proof. Suppose G is an acyclic digraph with $a \rightharpoonup b \rightharpoonup c$ as an induced subgraph, meaning here that a and c share no edge. It follows that the subgraph $\{a, c\}$ is discrete and therefore has an automorphism swapping a and c. But this finite automorphism cannot extend to an automorphism π of the whole graph, since this would cause a directed cycle: $a \rightharpoonup b \rightharpoonup c = \pi(a) \rightharpoonup \pi(b) \rightharpoonup \pi(c) = a$, a contradiction. □

Essentially the same point is made by observing that the collection of acyclic digraphs does not enjoy the amalgamation property, since the discrete graph with two vertices embeds into $a \rightharpoonup b \rightharpoonup c$ in two ways, but these cannot be amalgamated in any acyclic graph, since any amalgamation would give rise to a directed cycle.

The purpose of using graded digraphs is to overcome the obstacle posed by observation 6 and enable the theory of universality and homogeneity in the context of acyclic digraphs. The situation here is that the countable random $\mathbb{Q}$-graded digraph is homogeneous as an $\mathbb{Q}$-graded digraph, where the automorphisms respect the edge relation and the values of the $\mathbb{Q}$-grading, but not as an unlabeled digraph, where the automorphism need only respect the edge relation. Note that for $a \rightharpoonup b \rightharpoonup c$ in the case of an $\mathbb{Q}$-graded graph, the value of a must be strictly smaller than the value of c, which exactly prevents us in swapping a and c when we must preserve the values. Meanwhile, if one forgets both the values and the direction of edges, then in fact the undirected graph underlying the countable random $\mathbb{Q}$-graded digraph is the same as the countable random graph.

Observation 7. *The graph obtained by ignoring the values and direction of the edges in the countable random $\mathbb{Q}$-graded digraph is isomorphic to the countable random graph.*

Proof. It suffices to show that the graph underlying the countable random $\mathbb{Q}$-graded digraph Γ exhibits the finite-pattern property characterizing the countable random graph, namely, that for any two disjoint sets A and C, there is a node connected to every edge in A and to no edge in C. Such a node exists in Γ by the universality and homogeneity of Γ: simply select a rational value not used by any node in A and observe that there is a larger graded graph realizing the desired pattern with respect to A and C and having that value; by universality there is a copy of this digraph inside Γ; by homogeneity, we may translate this copy to align with the given nodes of A and C. So the undirected ungraded digraph underlying Γ has the finite-pattern property, and consequently is isomorphic to the countable random graph, since the countable random graph has this property and a simple back-and-forth argument shows that any two countable graphs with this property are isomorphic. □

3. A strengthening of Ressayre's theorem

Every model of arithmetic comes along with a model of finite set theory via the Ackerman coding, obtained by defining the following relation on the natural numbers:

$$n\ E\ m \quad\Leftrightarrow\quad \text{the } n^{\text{th}} \text{ binary bit of } m \text{ is } 1.$$

In the standard model $\mathbb{N}$, this relation is well-founded and extensional, and it is an often-assigned elementary exercise to prove that the Mostowski collapse of E is precisely the hereditary finite sets.

$$\langle \mathbb{N}, E\rangle \quad\cong\quad \langle \mathrm{HF}, \in\rangle$$

The relation E can be defined inside any model of arithmetic M, leading to the corresponding model of finite set theory $\langle \mathrm{HF}^M, \in^M\rangle$, and conversely, the natural numbers of HF^M are isomorphic to M again. In this way, the models of arithmetic are in a natural one-to-one correspondence with the models of finite set theory (properly understood), and the two theories are bi-interpretable. Specifically, [KW07] shows that the models of PA give rise to all the models of the finite set theory $\mathrm{ZFC}^{\neg\infty} = \mathrm{ZFC} - \mathrm{Inf} + \neg\mathrm{Inf} + \mathrm{TC}$, where TC is the assertion that every set has a transitive closure, and conversely the natural numbers of any such model satisfies PA, in inverse fashion, making for a bi-interpretation of these two theories. The curious anomaly here is that while the axiom TC is provable and may be omitted if one replaces the usual foundation axiom in ZFC with the axiom scheme of $\in$-induction, nevertheless results in [ESV11] show, quite interestingly, that TC is not provable if one uses only the usual foundation axiom.

The theme of this article begins in earnest with the following remarkable theorem of Ressayre's.

Theorem 8 (Ressayre [Res86])**.** *If M is any nonstandard model of* PA*, with $\langle \mathrm{HF}^M, \in^M\rangle$ the corresponding nonstandard hereditary finite sets of M, then for any consistent computably axiomatized theory T extending* ZF *in the language of set theory, there is a submodel $N \subseteq \langle \mathrm{HF}^M, \in^M\rangle$ such that $N \models T$.*

In particular, we may find models of ZFC and even of ZFC + large cardinals as submodels of HF^M, a structure whose theory is all about the finite and which thinks every object is finite. But nevertheless inside HF^M there are submodels of the strongest infinitary theories we know! Incredible! How can this be? Ressayre's proof uses partial saturation and resplendency to prove that one can find the submodel of the desired theory T.

Theorem 9 strengthens Ressayre's theorem, while simplifying the proof, by replacing the use of resplendency with a soft appeal to digraph universality. In particular, the role of the theory T is omitted: we need not assume that T is computable, and we don't just get one model of T, but rather all countable models of T, for the theorem shows that the nonstandard models of finite set theory are universal for all countable acyclic binary relations. In particular, every model of set theory arises as a submodel of $\langle \mathrm{HF}^M, \in^M\rangle$.

Theorem 9 (Main theorem 3)**.** *If M is any nonstandard model of* PA *or indeed merely of IE_1, then every countable model of set theory is isomorphic to a submodel of the hereditarily finite sets $\langle \mathrm{HF}^M, \in^M \rangle$ of M. Indeed, $\langle \mathrm{HF}^M, \in^M \rangle$ is universal for all countable acyclic binary relations.*

One may equivalently cast the theorem entirely in terms of finite set theory: every nonstandard model $\langle M, \in^M \rangle$ of finite set theory $\mathrm{ZFC}^{\neg\infty}$ is universal for all countable acyclic binary relations. In particular, every model of set theory, of ZFC or whatever theory, is isomorphic to a submodel of $\langle M, \in^M \rangle$, even though this latter model believes that every set is finite.

In order to prove theorem 9, we begin with a simple link between acyclic digraphs and sets with the set-membership relation.

Lemma 10. *Every finite acyclic digraph is isomorphic to a hereditarily finite set with the set-membership relation $\in$.*

Proof. Suppose that G is a finite acyclic digraph with n vertices $v_1, \dots v_n$. We first extend G to a larger finite extensional digraph G^+ as follows. Let N be the transitive closure of the digraph

$$0 \rightharpoonup 1 \rightharpoonup \cdots \rightharpoonup n,$$

which is the same as $(n+1, \in)$, if one uses the von Neumann ordinals as natural numbers, and let G^+ be the graph containing the disjoint union of G and N, together with edges $k \rightharpoonup v_k$ for $k \geq 1$. Note that (i) any two nodes of G have different predecessors in N; (ii) different nodes of N have different predecessors in N; and (iii) every node in N has 0 as a predecessor, but no node in G has 0 as a predecessor. Thus, no two nodes of G^+ have the same predecessors and so G^+ is extensional. Furthermore, since G and N are both acyclic and we added only edges pointing from N to G, it follows that G^+ is acyclic. In other words, the edge relation of G^+ is an extensional well-founded relation and therefore by the Mostowski collapse $\pi(v) = \{\, \pi(w) \mid w \rightharpoonup v \,\}$ is isomorphic to a unique transitive set under the set membership relation $\in$.

$$w \rightharpoonup v \qquad \text{if and only if} \qquad \pi(w) \in \pi(v).$$

This isomorphism associates each node in G^+ with the set of sets associated with its children. The isomorphism π carries the directed graph $(G, \rightharpoonup)$ to the hereditarily finite set $(A, \in)$, where $A = \{\, \pi(v) \mid v \in G \,\}$, as desired. □

Note that this lemma is provable in finite set theory $\mathrm{ZFC}^{\neg\infty}$. It will be generalized by lemma 21, which handles the general case of well-founded digraphs.

Proof of theorem 9. Suppose that M is a nonstandard model of PA, and that $\langle \mathrm{HF}^M, \in^M \rangle$ is the corresponding nonstandard model of finite set theory $\mathrm{ZFC}^{\neg\infty}$ built via the Ackerman coding in M. Consider the computable construction of the countable random $\mathbb{Q}$-graded digraph Γ, as carried out inside M. Let k be a nonstandard integer of M, and consider the k^{th} stage of this construction Γ_k^M. Since the standard stages of construction will be the same inside M as externally, it follows that the standard part of Γ_k^M is the actual countable random $\mathbb{Q}$-graded digraph Γ itself. That is, Γ is an induced subgraph of Γ_k^M. Since M thinks Γ_k is finite, we may apply lemma 10 inside M to find a set $A \in \mathrm{HF}^M$ such that M thinks $\langle \Gamma_k, \rightharpoonup \rangle \cong \langle A, \in^M \rangle$. Since Γ is an induced subgraph of Γ_k^M, we may restrict this isomorphism (externally) to find $\langle \Gamma, \rightharpoonup \rangle \cong \langle B, \in^M \rangle$ for some $B \subseteq \{\, a \in \mathrm{HF}^M \mid M \models a \in A \,\}$. Thus, we have realized the countable random $\mathbb{Q}$-graded digraph as isomorphic to a collection of sets in HF^M. But since Γ is universal among all countable acyclic digraphs, it follows that every countable set with an acyclic binary relation is isomorphic to an induced subgraph of Γ, and hence to a subcollection of HF^M under $\in^M$, as desired.

Lastly, let me discuss the issue of weakening the theory PA to $I\Delta_0$, which allows induction only for Δ_0 assertions, and even down to IE_1, which has induction only for bounded existential assertions. The reason is that every countable nonstandard model of IE_1 has an initial segment that is a model of PA, by a result of Paris [Par84], whose proof relies on the earlier corresponding result for $I\Delta_0$, proved by McAloon [McA82]. So if we begin with a model $M \models IE_1$, therefore, we may simply cut down to the PA initial segment $M_0 \models \mathrm{PA}$, where we will find the universal structure by the argument of the previous paragraph. Thus, every model of an acyclic binary relation embeds into $\langle \mathrm{HF}^{M_0}, \in^{M_0} \rangle$, which is an initial segment of $\langle \mathrm{HF}^M, \in^M \rangle$, and so every model of an acyclic binary relation embeds into HF^M, as desired. □

Corollary 11. *The countable models of set theory, up to isomorphism, have a universal object under the submodel relation. Namely, if $\langle M, \in^M \rangle$ is any ω-nonstandard model of set theory, then every countable model of set theory is isomorphic to a submodel of $\langle M, \in^M \rangle$.*

Proof. Indeed, theorem 9 shows that every countable model $\langle N, E \rangle$ of any binary relation embeds already into the hereditary finite sets $\langle \mathrm{HF}^M, \in^M \rangle$ of M. □

I will show later that in fact every countable nonstandard model of set theory is universal in this way, even if the ill-foundedness appears only higher than ω.

4. ℓ-graded digraphs and the finite-pattern property

It will be convenient to consider grading notions other than $\mathbb{Q}$, aiming eventually at the well-founded acylic digraphs, which are precisely the α-graded digraphs for some ordinal α, or in the case of proper classes, the Ord-graded digraphs, as well as the No-graded digraphs, using the surreal numbers No. To assist with this analysis, consider the general case of an arbitrary linear order ℓ and the ℓ-graded digraphs G, which have a value assignment $v \mapsto \alpha_v$ of the nodes $v \in G$ with elements $\alpha_v \in \ell$ in such a way that $v \rightharpoonup w$ in G implies $\alpha_v < \alpha_w$ in ℓ. Any such graph is acyclic, since a directed cycle in G would give rise to a violation of asymmetry in ℓ.

Theorem 12. *For any countable linear order ℓ, there is a countable homogeneous ℓ-graded digraph that is universal for all countable ℓ-graded digraphs. This digraph is unique up to isomorphism (even when restricting universality to the finite ℓ-graded digraphs), and has a computable presentation if ℓ is computable.*

Proof. The collection $\mathcal{Q}_\ell$ of all finite ℓ-graded directed graphs has the hereditary property, the joint embedding property and the amalgamation property, and consequently admits a Fraïssé limit, which is up to isomorphism the unique countable, homogeneous, ℓ-graded digraph of its age. This graph is universal for all countable ℓ-graded directed graphs by the same argument as in theorem 5.

Alternatively, it may be easier simply to view ℓ as a suborder of the rational order $\mathbb{Q}$, and then consider the restriction $\Gamma \restriction \ell$ of the countable random $\mathbb{Q}$-graded digraph Γ to the vertices having a label in ℓ. The subgraph $\Gamma \upharpoonright \ell$ is ℓ-graded and inherits the homogeneity property from Γ, since any finite partial automorphism of $\Gamma \upharpoonright \ell$ is also a finite partial automorphism of Γ, which therefore extends to an automorphism of Γ, which respects the labels and thus provides an automorphism of $\Gamma \upharpoonright \ell$. The digraph $\Gamma \upharpoonright \ell$ remains universal for ℓ-graded directed graphs, because any such graph embeds into Γ in a way that uses only the labels in ℓ and hence embeds into $\Gamma \upharpoonright \ell$.

If ℓ is computable, then ℓ has a computable copy inside $\mathbb{Q}$, and using the computable presentation of Γ from theorem 5, we obtain a computable presentation $\Gamma \upharpoonright \ell$, as desired. □

Let us call this graph the *countable random ℓ-graded digraph.* There is also a finite-pattern property for these graphs analogous to the ω-categorical first-order characterization of the countable random graph, mentioned in observation 7.

Definition 13. An ℓ-graded digraph Γ satisfies the *finite-pattern property*, if for any disjoint finite sets of vertices A, B and C and any $\alpha \in \ell$, such that every vertex in A has value less than α and every vertex in B has value greater than α, then there is a vertex v in Γ with value exactly α such that $a \rightharpoonup v$ and $v \rightharpoonup b$ for every $a \in A$ and $b \in B$, but v has no edges with any vertex in C.

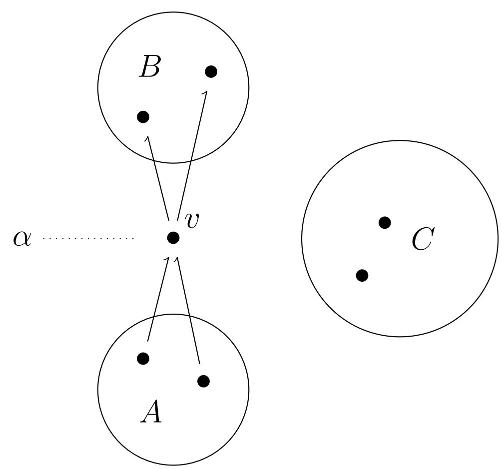

FIGURE 1. The finite-pattern property

Equivalently, an ℓ-graded digraph has the finite-pattern property just in case it is existentially closed in the class of all ℓ-graded digraphs.

Theorem 14. *If Γ is a countable ℓ-graded digraph for a countable linear order ℓ, then the following are equivalent:*

(1) Γ *has the finite-pattern property.*
(2) Γ *is isomorphic to the countable random ℓ-graded digraph.*

Proof. It is clear that the countable random ℓ-graded digraph has the finite-pattern property, since each instance with particular A, B and C is clearly realizable in a finite ℓ-graded digraph, which must therefore embed into the countable random ℓ-graded digraph by universality, thereby providing a realizing node v by homogeneity.

Conversely, suppose that Γ is a countable ℓ-graded digraph with the finite-pattern property. It is clear by successive applications of this property that every ℓ-graded directed graph occurs as an induced subgraph of Γ. Now, one simply applies a back-and-forth construction to realize that Γ is isomorphic to the countable random ℓ-graded digraph. At each stage, one applies the finite-pattern property either in Γ or in the countable random ℓ-graded digraph to extend the isomorphism one more step. □

Although the random digraphs that we have constructed are homogenous as $\mathbb{Q}$-graded or as ℓ-graded digraphs, as in observation 6 they are definitely not homogeneous merely as digraphs, without the extra structure provided by the value assignment. Meanwhile, if one should forget the values and the directionality of the edges, then the argument of observation 7 shows that the graph underlying the countable random ℓ-graded digraph, for any infinite linear order ℓ, is isomorphic to the countable random graph.

5. Uncountable homogeneous graphs and structures

The main argument of subsequent sections will make use of uncountable homogeneous graded digraphs, and it will be convenient to have a general theory of how to construct such structures. So let us briefly review that theory here, beginning with the uncountable analogue of the Fraïssé limit construction.

Although the Fraïssé limit construction is typically given for the countable case, where one produces a countable homogeneous structure as a limit of finitely generated structures, the idea generalizes naturally to uncountable structures, and I shall give a brief account here of how this can be done.

For any first-order structure D, the *age* of D is the class of finitely generated structures that embed into D, that is, which are isomorphic to a substructure of D, and for any cardinal δ, the δ-*age* is the class of $<\delta$-generated structures that embed into D. A structure D is δ-*homogeneous* if every isomorphism of two $<\delta$-generated substructures of D extends to an automorphism of D.

Theorem 15. *Suppose that K is a family of structures in a language $\mathcal{L}$ of size at most δ, an infinite cardinal for which $\delta^{<\delta} = \delta$, such that K is closed under isomorphism, contains at most δ many isomorphism types, and every member of K is finitely generated. Suppose further that K exhibits the the following properties:*

- HP *(hereditary property) Every submodel of a model in K is in K.*
- JEP *(joint embedding property) If $A, B \in K$ there is $C \in K$ with A and B both embedding into C.*
- AP *(amalgamation property) For any $A, B, C \in K$, if $e : A \to B$ and $f : A \to C$ are embeddings, then there is $D \in K$ with embeddings $g : B \to D$ and $h : C \to D$ such that $ge = hf$.*

Then there is a δ-homogeneous structure $\mathcal{A}$ of size δ whose age is exactly K. This structure is unique up to isomorphism and universal among all structures of size δ whose age is contained in K.

Proof. See Hodges [Hod93, thm 7.1.2] for an account of the countable case $\delta = \omega$. Assume the class K has HP, JEP and AP. Let K^+ be the class of $<\delta$-generated structures whose age is contained in K. Every element of K^+ is the direct limit of its finitely generated submodels, and these are all in K. Thus, K^+ is simply the closure of K under limits of directed systems of size less than δ. Clearly, K^+ retains the HP. In fact, K^+ also enjoys the JEP and the AP, which we prove jointly by induction on the number of generators for the structure involved. For the JEP, consider two structures $A, B \in K^+$, where A is γ-generated and B is ξ-generated, where $\xi \leq \gamma$. The case of finite γ follows from the JEP of K. Otherwise, assume γ is an infinite cardinal and express A as the union of a chain $A = \bigcup_{\alpha<\gamma} A_\alpha$ of $<\gamma$-generated structures. By the AP induction hypothesis applied to A_α and B, we may construct a chain $C_0 \subseteq \cdots \subseteq C_\alpha \subseteq \cdots$ for $\alpha < \gamma$, such that A_α and B embed into C_α, and the embedding of A_α into C_α extends to the embedding of $A_{\alpha'}$ into $C_{\alpha'}$ for $\alpha < \alpha'$. The union $C = \bigcup_\alpha C_\alpha$ jointly embeds both A and B. A similar argument establishes the AP for K^+. In addition to HP, JEP and AP, the class K^+ also enjoys the $<\delta$-*limit property* ($<\delta$-LP), which asserts that if each $A_\alpha \in K^+$ and

$$A_0 \subseteq A_1 \subseteq \cdots \subseteq A_\alpha \subseteq \cdots$$

is a chain of length less than δ, then $\bigcup_\alpha A_\alpha \in K^+$.

Now, in a forcing-style presentation, let $\mathbb{P}$ be the structures in K^+ whose underlying set is contained in δ. (In the case that the language has δ many constant symbols, then let us also fix a bijection $\delta \cong \delta \times \delta$ and insist furthermore that the β-generated structures of $\mathbb{P}$ live in the part of δ corresponding to $\beta \times \delta$; this will ensure that there is sufficient extra room to expand the structure.) Every structure in K^+ is isomorphic to an element in $\mathbb{P}$. We order $\mathbb{P}$ by the submodel relation, and consider it as a forcing notion, for which the larger models are stronger. Note that $\mathbb{P}$ is closed under the unions of chains of size less than δ. The fact that K^+ has the JEP means that for any $A \in K^+$ the collection of conditions $B \in \mathbb{P}$ such that A embeds into B is dense in $\mathbb{P}$. The fact that K^+ has the AP means that for any $A, B, C \in \mathbb{P}$ such that $A \subseteq B$ and $f : A \to C$, the set of conditions $D \in \mathbb{P}$ such that $C \subseteq D \Rightarrow f$ extends to some $g : B \to D$, is dense below C in $\mathbb{P}$. Since $\delta^{<\delta} = \delta$, there are only δ many conditions in $\mathbb{P}$ and thus only δ many dense sets of the type I have just mentioned in the previous sentences. Since $\mathbb{P}$ is also closed under the unions of chains of length less than δ, we may by build a length δ chain of structures

$$D_0 \subseteq \cdots \subseteq D_\alpha \subseteq \cdots$$

from $\mathbb{P}$ in such a way so as to meet all these dense sets. It follows that the union $D = \bigcup_{\alpha<\delta} D_\alpha$ has size δ and contains isomorphic copies of every $A \in K^+$. Further, the second family of dense sets exactly ensure that the structure D is δ-*weakly homogeneous*, meaning that whenever $A \subseteq B$ and $f : A \to D$ is an embedding, where $A, B \in K^+$, then f extends to some $g : B \to D$. It follows that D is universal for any structure B of size δ, whose age is contained in K. To see this, first express $B = \bigcup_{\alpha<\delta} B_\alpha$ as a union of structures $B_\alpha \in K^+$, and then build a coherent system of embeddings $f_\alpha : B_\alpha \to D$, applying the δ-weak homogeneity property of D to extend at successor stages. The result is an embedding of B into D, as desired. A similar back-and-forth argument shows that D is δ-homogeneous, and uniqueness follows from this. □

Note that the $<\delta$-limit property may be equivalently formulated as the assertion that the class is closed under direct limits of linear systems from K of length less than δ.

A class structure is *set-homogeneous* if every isomorphism between two set substructures extends to a class automorphism of the whole structure. The *global choice* principle asserts that there is a proper class choice function selecting an element from every set. This is equivalent to the assertion that there is a proper class well-ordering of the universe in order type Ord. The axiom is most naturally stated in the second-order Gödel-Bernays set theory, which allows for classes that may not be definable, and so the global choice axiom does not necessarily imply $V = \text{HOD}$.

Theorem 16. *Assume the global choice principle. Suppose that K is a proper class of finitely generated models in language $\mathcal{L}$, which may be a proper class, and that K is closed under isomorphic copies and exhibits the HP, JEP and AP. Then there is an set-homogeneous class model $\mathcal{A}$, unique up to isomorphism, whose age is exactly K. Furthermore, every class model whose age is contained in K is embeddable into $\mathcal{A}$.*

Proof. This is simply a proper class version of theorem 15. One sets up the class partial order $\mathbb{P}$ consisting of all set-generated structures, built on the ordinals, whose age is contained in K, and order them by induced submodel as before. The point is that this class forcing is set-closed, just as the forcing notion in theorem 15 was $<\delta$-closed, and so using the global choice principle, one may construct an Ord-length sequence of structures that meet all the dense classes corresponding to those in the proof of theorem 15. The result is a class structure $\mathcal{A}$ that is set-homogeneous and universal for all class structures whose age is contained in K. □

I would like now to construct canonical homogeneous ℓ-graded digraphs for an arbitrary linear order ℓ, perhaps uncountable. Generalizing definition 13 to the uncountable, we say for a cardinal δ that an ℓ-graded digraph G exhibits the *$<\delta$-pattern property* if for any disjoint sets A, B and C of nodes, with cardinality less than δ, and any $\alpha \in \ell$ for which every node in A has value less than α and every node in B has value greater than α, there is a node $v \in \Lambda$ with value exactly α such that $a \rightharpoonup v$ and $v \rightharpoonup b$ for all $a \in A$ and $b \in B$, while also $v \perp c$ for all $c \in C$.

As a first-order structure, we shall regard an ℓ-graded digraph as a structure $\langle G, \rightharpoonup, U_\alpha \rangle_{\alpha \in \ell}$, where $\langle G, \rightharpoonup \rangle$ is the underlying digraph and U_α is a unary relation holding of the vertices with value α.

Theorem 17. *For any linear order ℓ and any cardinal $\delta \geq |\ell|$ for which $\delta^{<\delta} = \delta$, there is a δ-homogenous ℓ-graded digraph of size δ, which is universal for ℓ-graded digraphs of size at most δ. This graph is unique up to isomorphism and satisfies $<\delta$-pattern property.*

Proof. Let $\mathcal{Q}$ be the collection of all finite ℓ-graded digraphs. This collection has the hereditary property, the joint embedding property and the amalgamation property. Thus, by theorem 15 it has a δ-Fraïssé limit Λ, which is a δ-homogenous ℓ-graded digraph of size δ, universal for the ℓ-graded digraphs in $\mathcal{Q}$, and unique up to isomorphism with these properties. Further, it is universal among all ℓ-graded digraphs of size at most δ. From this and δ-homogeneity, it follows that Λ has the $<\delta$-pattern property. □

Theorem 18. *If the global axiom of choice holds, then there is a set-homogeneous* Ord*-graded class digraph, which is universal for all* Ord*-graded class digraphs.*

Proof. Appeal to theorem 16. □

The remarks after theorem 20 show that this digraph admits a canonical construction in the manner of the surreal numbers.

6. The surreal numbers and the hypnagogic digraph

In theorem 18 we constructed a set-homogeneous Ord-graded class digraph that was universal for all Ord-graded digraphs. This digraph arises as the class Fraïssé limit of all Ord-graded finite digraphs, and it is unique up to isomorphism. In this section, I would like to describe how this graph also arises by a canonical construction intimately connected with the construction of the surreal numbers and what I call the hypnagogic digraph.

Let us first recall the standard construction of the class of surreal numbers, denoted No. Here, I am concerned with the surreal numbers only as a linear order; the ordered field structure and other structure on the surreals will not be relevant here.[1] To construct the surreal number line, one begins with nothing and then proceeds relentlessly, transfinitely, to fill all possible cuts in the order created so far. One defines the class No of all surreal numbers and their order in a simultaneous recursion, for which a surreal number is born at an ordinal stage if it is represented by a pair $\{\,A \mid B\,\}$, where A and B are sets of previously-born surreal numbers, where every number in A is smaller than every number in B. The idea is that $\{\,A \mid B\,\}$ will be a surreal number filling the cut between A and B, so that it is larger than every element of A and smaller than every element of B. The very first surreal number to be created is $0 = \{\ \mid\ \}$, meaning by this notation that we have the empty set on each side, and then immediately afterwards the cut above $1 = \{\,0 \mid\ \}$, meaning the singleton set on the left and the empty set on the right, and similarly the cut below $-1 = \{\ \mid 0\,\}$. Just as with the rational numbers, however, the left-set/right-set $\{\,A \mid B\,\}$ representation of the surreal numbers will not be unique, for different pairs of sets can determine the same cut, and so we will quotient by an underlying equivalence relation. Specifically, for surreal numbers $x = \{\,X_L \mid X_R\,\}$ and $y = \{\,Y_L \mid Y_R\,\}$, we define $x \leq y$ if and only if no obvious obstacle prevents it, namely, there is no $x_L \in X_L$ with $y \leq x_L$ and no $y_R \in Y_R$ with $y_R \leq x$. Two numbers x and y are equivalent when $x \leq y \leq x$, and in this case we say that x and y are equal as surreal numbers. It is easy to prove that a surreal number $\{\,A \mid B\,\}$ is equivalent to the first-born surreal number that is bigger than every element of A and smaller than every element of B.

One sees immediately the similarity in the construction of the surreal numbers and the construction of the proper class Fraïssé limit of theorem 16. And indeed, by those methods one can see that the surreal numbers are simply the proper class Fraïssé limit of the class of all finite linear orders. They are set-homogeneous and universal for all proper class linear orders (provided global choice holds, but one can construct them and get universality for set linear orders just in ZFC). In this sense, the surreal number line is for proper class linear orders what the rational number line $\mathbb{Q}$ is for the countable linear orders. What I aim to do in this section is to continue this analogy to find the corresponding proper class analogue of the countable random $\mathbb{Q}$-graded digraph. The answer is what I call the *hypnagogic digraph*, denoted Hg. To explain the terminology, the hypnagogic state is the dream-like, sometimes hallucinatory state at the boundary between wakefulness and sleep, and its usage here is meant to evoke both the universal property of the graph as well as its homogeneity: peering

[1] Considerably pre-dating the surreal numbers, the study of what we now call the saturated linear orders was undertaken first by Hausdorff [Hau02], who introduced the notation η for the countable dense linear order and η_α for the saturated dense linear order of size $\aleph_\alpha$, when $\aleph_\alpha^{<\aleph_\alpha} = \aleph_\alpha$. In this notation, the surreal number line is simply η_{Ord}.

into the hypnagogic digraph, one may see all the various dancing visions of any given graded digraph.

Theorem 19. *There is surreal-numbers-graded class digraph* Hg*, the hypnagogic digraph, such that:*

1. Hg *is universal for all graded digraphs using any linear order, as a suborder of* No*.*
2. Hg *is set-homogeneous, meaning that any isomorphism of two induced graded subgraphs extends to an automorphism of* Hg*.*
3. *If global choice holds, then* Hg *is universal for all proper class graded digraphs, using any class linear order as a suborder of* No*, and furthermore is uniquely determined up to isomorphism by (1) and (2).*
4. Hg *exhibits the set-pattern property.*

Proof. Just as for the countable random $\mathbb{Q}$-graded digraph in theorem 5, there are numerous constructions of this highly canonical object. For example, the hypnagogic digraph is the Fraïssé limit of the class of all finite No-graded digraphs, as described in theorem 16. It is simply the Ord-saturated No-graded class digraph. But meanwhile, there is also a highly canonical construction of the hypnagogic digraph, closely connected with the construction of the surreal numbers, which I prefer to take as the official definition of the class Hg. The idea is simply to use the left-set/right-set representation $\{ A \mid B \}$ of the surreal numbers, *but do not quotient by the equivalence relation!* Rather, the hypnagogic digraph is the term algebra of the terms representing surreal numbers. That is, the nodes in Hg are all the various representations $\{ A \mid B \}$ of surreal numbers, counted as different nodes even if they are equal as surreal numbers. For a given node $v = \{ A \mid B \}$, we place edges $a \rightharpoonup v$ and $v \rightharpoonup b$ for every $a \in A$ and $b \in B$, and we define the value of v to be simply the surreal number that it represents. Thus, the meaning of $\{ A \mid B \}$ as a node in Hg is that it is pointed at from every $a \in A$; it points at every $b \in B$; and its value is $\{ A \mid B \}$ as a surreal number. Taking a big-picture view of the construction, the hypnagogic digraph is the result of systematically adding nodes with a specified edge connectivity with the already constructed nodes. Under this description, it is clear that Hg satisfies the set-pattern property, which in turn implies the universality property and thus the set-homogeneity by a class-length back and forth property. (Note: one need not use the global choice principle for the set-homogeneity of Hg, since any automorphism of the nodes up to a given birthday induces a canonical automorphism of the cuts created at that birthday.) One gets universality for all class ℓ-graded digraphs, using any linear order ℓ, simply because under the global choice principle every linear order embeds as a suborder of No, and then one can systematically build an embedding of the given ℓ-graded digraph by mapping to a node representing the desired value and connectivity. Uniqueness follows similarly by a class-length back-and-forth construction. □

One way to view the construction of the hypnagogic digraph is that when a surreal number $v = \{ A \mid B \}$ is born, the meaning of this term is that v is definitely larger than every element of A and smaller than every element of B. The hypnagogic digraph takes this meaning seriously by adding digraph edges exactly to record this information. In this way, the hypnagogic digraph seems fundamentally to underlie the surreal numbers, in a sense existing just prior to the surreal numbers.

To further support the idea that the hypnagogic digraph Hg is the proper class analogue of the countable random $\mathbb{Q}$-graded digraph, observe that the hypnagogic digraph Hg^M as computed inside a countable model $M \models \mathrm{ZFC}$ is externally isomorphic to the countable random $\mathbb{Q}$-graded digraph. This is because the surreal numbers No^M as computed in M are a countable dense linear order without endpoints, and thus isomorphic to the rational line $\mathbb{Q}$, and Hg^M is thus essentially a countable $\mathbb{Q}$-graded digraph with the finite-pattern property, which is isomorphic to the countable random $\mathbb{Q}$-graded digraph by theorem 14.

The surreal number line No is universal for all class linear orders. For any such suborder $\ell \subseteq \mathrm{No}$, we may form the *hypnagogic ℓ-graded digraph* $\mathrm{Hg} \restriction \ell$, simply by restricting the hypnagogic digraph to the nodes with value in ℓ.

Theorem 20. *For any class linear order $\ell \subseteq \mathrm{No}$, the hypnagogic ℓ-graded digraph $\mathrm{Hg} \restriction \ell$ is set-homogeneous and universal for all ℓ-graded digraphs (including class digraphs if global choice holds). Furthermore, it is uniquely determined by these properties.*

Proof. Clearly $\mathrm{Hg} \restriction \ell$ is an ℓ-graded digraph, and it inherits the set-homogeneity of Hg, since any isomorphism of induced subgraphs of $\mathrm{Hg} \restriction \ell$ is also an isomorphism of induced subgraphs of Hg, which therefore extends to an automorphism of all of Hg, which provides an automorphism of $\mathrm{Hg} \restriction \ell$ since these automorphisms respect the values of nodes. Furthermore, $\mathrm{Hg} \restriction \ell$ is universal for all ℓ-graded digraphs, since Hg is, by embeddings that respect the values of nodes. Uniqueness follows (assuming the global choice principle) by the usual back-and-forth argument. □

In particular, since the ordinals are a suborder of the surreal numbers, $\mathrm{Ord} \subseteq \mathrm{No}$, we have the hypnagogic Ord-graded digraph $\mathrm{Hg} \restriction \mathrm{Ord}$, which by the uniqueness claim of theorem 20 is the same Ord-graded digraph constructed in theorem 18. One can view the hypnagogic Ord-graded digraph as constructed in a grand recursion: at each stage, for any ordinal α up to that stage and any sets A and B of previously constructed nodes, with those in A having value less than α and those in B above, one forms a new node v with value α and creates edges $a \rightharpoonup v$ and $v \rightharpoonup b$ for every $a \in A$ and $b \in B$.

Note that if M is a countable transitive model of set theory, with height $\lambda = \mathrm{Ord}^M$, then the hypnagogic Ord-graded digraph $(\mathrm{Hg} \restriction \mathrm{Ord})^M$ as constructed in M is isomorphic by theorem 14 to the countable random λ-graded digraph.

7. Realizing well-founded digraphs as sets

The following basic lemma, a modification of the Mostowski collapse construction, shows how to realize well-founded acyclic digraphs as sets. Similar modified Mostowski collapse maps have arisen for diverse purposes, such as in the non-well-founded set theory of Aczel [Acz88] and in the finite set theory of Kirby [Kir10] and elsewhere.

Lemma 21. *Every well-founded digraph $(G, \rightharpoonup)$ is isomorphic to a set $(A, \in)$ under the set-membership relation $\in$. Furthermore,*

(1) *If G is finite, then there is a hereditarily finite such set A.*
(2) *If G is countable and λ-graded for some ordinal λ, then A can be found with $A \subseteq V_{\omega+\lambda}$. In particular, if $\omega^2 \leq \lambda$, then there is such $A \subseteq V_\lambda$.*
(3) *If G is λ-graded and $|G| \leq |V_\beta|$, then A can be found with $A \subseteq V_{\beta+2+\lambda}$. In particular, if $\beta^2 \leq \lambda$, then there is such $A \subseteq V_\lambda$.*
(4) *If G is countable and λ-graded, with λ infinite, then A can be found with $A \subseteq L_{\lambda+\lambda}$.*

Proof. Suppose that $(G, \rightharpoonup)$ is a well-founded digraph. Define

$$\pi(x) = \{\, \pi(y) \mid y \rightharpoonup x \,\} \cup \{\, \{\, \emptyset, x \,\} \,\},$$

and let $A = \{\, \pi(x) \mid x \in G \,\}$. This definition is well-defined by recursion on the well-founded relation $\rightharpoonup$, since the value of $\pi(x)$ is determined by the values of $\pi(y)$ for earlier $y \rightharpoonup x$. I claim that it is an isomorphism of $(G, \rightharpoonup)$ with $(A, \in)$. It is essentially similar, of course, to the Mostowski collapse, but modified by the inclusion of the extra element $\{\, \emptyset, x \,\}$, which is added in order to distinguish the nodes sufficiently. (In the case that the original graph is extensional, then the nodes are already distinguished by their predecessors and this modification is unnecessary, for the Mostowski collapse is already an isomorphism.) It is clear from the definition that $y \rightharpoonup x$ implies $\pi(y) \in \pi(x)$. For the converse direction, note first that no $\pi(x)$ is empty and since $\{\, \emptyset, x \,\}$ is nonempty, it follows also that $\emptyset \notin \pi(y)$ for any y. In particular, $\pi(y) \neq \{\, \emptyset, x \,\}$ for any x and y. Thus, if $\pi(y) \in \pi(x)$, it must be that $\pi(y) = \pi(z)$ for some $z \rightharpoonup x$. Since $\{\, \emptyset, y \,\} \in \pi(y)$ and $\{\, \emptyset, z \,\} \in \pi(z)$, but these sets are not $\pi(u)$ for any u, it follows that $\{\, \emptyset, y \,\} = \{\, \emptyset, z \,\}$ and consequently $y = z$ and thus actually $y \rightharpoonup x$. So I have established for all $x, y \in G$ that

$$y \rightharpoonup x \qquad \text{if and only if} \qquad \pi(y) \in \pi(x),$$

and this is precisely what it means for π to be an isomorphism of $(G, \rightharpoonup)$ with $(A, \in)$. This establishes the basic claim of the lemma.

I now consider the further claims by analyzing the nature of A. If G is finite, then we may assume that the underlying nodes of G are natural numbers (or some other hereditary finite sets), and in particular, the sets $\{\, \emptyset, x \,\}$ arising in the definition of π are all hereditarily finite. It follows inductively that every $\pi(x)$ is hereditarily finite, and so A is a finite set of hereditarily finite sets and thus hereditarily finite itself, establishing (1). Similarly, for (2), if G is countable, then we may again assume that the nodes of G are natural numbers and consider the isomorphism π as defined above. By induction on values in the grading, it follows for any node x with value α that $\pi(x) \in V_{\omega+\alpha}$, since $\{\, \emptyset, x \,\} \in V_\omega$ and $\pi(x)$ consists otherwise of $\pi(y)$, where y has some value $\beta < \alpha$ and hence by induction $\pi(y) \in V_{\omega+\beta}$. So $A \subseteq V_{\omega+\lambda}$. If $\omega^2 \leq \lambda$, then $\omega + \lambda = \lambda$, and so in this case $A \subseteq V_\lambda$.

A similar argument works in the case of (3). Assume that G is λ-graded and the nodes of G come from V_β. In this case, $\{\, \emptyset, x \,\} \in V_{\beta+1}$ at worst, and inductively $\pi(x) \in V_{\beta+2+\alpha}$ when x has value α. And so $A \subseteq V_{\beta+2+\lambda}$. If $\beta^2 \leq \lambda$, then $\beta + 2 + \lambda = \lambda$, and so we have $A \subseteq V_\lambda$, as desired.

Assertion (4) is subtle. As in theorem 20, let $\Lambda = (\mathrm{Hg} \upharpoonright \lambda)^{L_\lambda}$ be the hypnagogic λ-graded digraph as defined in L_λ. Since this graph is universal for all countable λ-graded digraphs, the original digraph G embeds into Λ, and so it will suffice just to handle Λ. This is a definable class in L_λ. Although the edge relation $\rightharpoonup$ of Λ is definable in L_λ, we will not be able to carry out the modified Mostowski collapse of Λ inside L_λ, because Λ has λ many nodes of each given value, and indeed, every node in Λ has λ many predecessors there. Thus, the sets arising at every step of the modified Mostowski collapse will have size λ, and consequently will not necessarily be elements of L_λ. Another way to say it is that the edge relation $\rightharpoonup$ of Λ is not set-like in L_λ, and this is why if we want to use this graph directly, we have to build sets on top of L_λ, stretching up to $L_{\lambda+\lambda}$. Specifically, let $\pi : \Lambda \to A$ be the modified Mostowski collapse as defined above. I claim that $A \subseteq L_{\lambda+\lambda}$, by arguing that any node $x \in \Lambda$ of value α has $\pi(x) \in L_{\lambda+1+\alpha+1}$. This is true if x has value 0, since $x \in L_\lambda$ and $\pi(x) = \{\{\emptyset, x\}\} \in L_{\lambda+2}$ at worst (although if λ is a limit ordinal, one achieves L_λ here). More generally, if the claim is true for all values below α and x has value α, then $\pi(x) = \{\pi(y) \mid y \rightharpoonup x\} \cup \{\{\emptyset, x\}\}$, where for $y \rightharpoonup x$ we have $\pi(y) \in L_{\lambda+1+\beta}$ for some $\beta < \alpha$, and so $\pi(x)$ is a subset of $L_{\lambda+1+\alpha}$, which is definable from $\Lambda, x \in L_\lambda$, and so $\pi(x) \in L_{\lambda+1+\alpha+1}$. Thus, altogether, $A \subseteq L_{\lambda+\lambda}$. So we have found a set $A \subseteq L_{\lambda+\lambda}$ such that $(A, \in)$ is the countable random λ-graded digraph. It now follows by the universality property of this digraph identified in theorem 12 that any countable λ-graded digraph is isomorphic to a subset of this particular A, establishing (4). □

It will follow as a consequence of the main theorem, theorem 29, that we may actually improve statements (2) and (4) of lemma 21 to the following assertion:

(5) If G is countable and λ-graded with λ infinite, then $\langle G, \rightharpoonup\rangle \cong \langle A, \in\rangle$ for some $A \subseteq L_\lambda$.

Indeed, this will be the key point of the proof of the main theorem. Meanwhile, let us record in the following observation the reason we are interested in graded digraphs, as far as models of set theory are concerned.

Observation 22. *Every model of set theory $\langle M, \in^M\rangle$ is an Ord^M-graded digraph, using von Neumann rank as values.*

Proof. Surely Ord^M is a linear order, and $y \in x$ implies that y has lower rank than x. □

8. Warming up to the main theorems

In this section I explain how lemma 21 can be used to prove several approximations to main theorem 1, which I give here as a warm-up, since the proof of main theorem 1 will introduce several complications.

Proposition 23. *If M is any countable transitive model of set theory and $\lambda = \mathrm{Ord}^M$, then $\langle M, \in\rangle$ is isomorphic to a submodel of $\langle L_{\lambda+\lambda}, \in\rangle$.*

Proof. This proposition is a quick corollary to statement (4) of lemma 21, since if we view $\langle M, \in^M\rangle$ as a λ-graded digraph as in observation 22, then lemma 21 statement (4) says that it is isomorphic to a submodel of $\langle L_{\lambda+\lambda}, \in\rangle$, as desired. □

The need to go to $L_{\lambda+\lambda}$ in proposition 23 and in lemma 21 statement (4) is directly connected with the fact that the hypnagogic λ-graded digraph Λ as computed in L_λ is not set-like in L_λ, since every node has λ many predecessors of each smaller value. Perhaps one might hope to overcome this difficulty by finding a λ-graded digraph in L_λ that was both set-like and had the finite-pattern property, for in this case, it would be universal for all λ-graded binary relations and we would be able to perform the modified Mostowski collapse of lemma 21 inside L_λ itself. This would prove that M is isomorphic to a submodel of $L_\lambda = L^M$, which is the main goal here. Unfortunately, however, the following result shows that there is no such digraph in any model of ZF.

Observation 24. *No model of* ZF *has an* Ord*-graded digraph class Γ with the finite-pattern property, such that Γ is set-like.*

Proof. Suppose that Γ is a set-like Ord-graded digraph class with the finite-pattern property. Consider any fixed node p with nonzero value, and let A_p be the set of predecessors of p in Γ, which is a set precisely because we assumed that Γ is set-like. Since A_p has only a set number of subsets, but there are a proper class of higher levels, it must be that there are two distinct nodes q, r with value above p, but with the same pattern of predecessors on A_p. That is, for $v \in A_p$, we have $v \rightharpoonup q$ if and only if $v \rightharpoonup r$. This violates the finite-pattern property, which would require the existence of a node v with $v \rightharpoonup p$ and $v \rightharpoonup q$ but $v \perp r$. Such a vertex v would reveal that q and r have different predecessors below p, contrary to assumption. □

The complications introduced into the proof of the main theorem, using the surrogate parent concept, are aimed specifically at overcoming this difficulty. Meanwhile, we can also overcome the difficulty by relaxing to a weaker but still commonly considered theory. Let "$V = H_{\kappa^+}$" be the theory asserting that κ is the largest cardinal and that every set has hereditary size at most κ. For any cardinal κ, the collection H_{κ^+} is a model of $\mathrm{ZFC}^- + V = H_{\kappa^+}$, and these structures and their elementary substructures are very commonly considered in set theory.

Theorem 25. *Suppose M is a countable transitive model of $\mathrm{ZFC}^- + V = H_{\kappa^+}$, where $\kappa^{<\kappa} = \kappa$ in M and $\lambda = \mathrm{Ord}^M$. Then there is a λ-graded digraph $\Lambda \subseteq M$, which is set-like in M and which obeys the $<\kappa$-pattern property in M.*

Proof. We modify the construction of the hypnagogic λ-graded digraph so that it will become set-like, by restricting it to have κ many nodes of any given value. Specifically, for each $\beta < \lambda$ of size κ, consider the construction in M of the β-graded hypnagogic digraph as the Fraïssé limit of the finite β-graded digraphs. This graph is an element of M and has the $<\kappa$-pattern property in M. Furthermore, for any $\beta < \gamma < \lambda$ and any copy of the β-graded hypnagogic digraph Λ_β, it may be extended to the hypnagogic γ-graded digraph Λ_γ as constructed in M, and furthermore extended in such a way that Λ_β is precisely $\Lambda_\gamma \upharpoonright \beta$. This is simply because $\Lambda_\gamma \upharpoonright \beta$ is κ-homogeneous and has the $<\kappa$-pattern property, and hence is isomorphic to Λ_β. Continuing the main argument, now, I select an increasing sequence $\langle \lambda_n \mid n < \omega \rangle$ cofinal in λ and form the limit graph $\Lambda = \bigcup_n \Lambda_{\lambda_n}$, where the Λ_n form a coherent sequence of induced subgraphs, with $\Lambda_{\lambda_n} = \Lambda \upharpoonright \lambda_n$. The graph Λ is set-like in M, because every node lives in some Λ_n, which is an element of M, and all its children are also in Λ_n. Furthermore, Λ is λ-graded and has the finite-pattern property. Moreover, each Λ_n and consequently also Λ has the $<\kappa$-pattern property with respect to M, since every subset of Λ that is in M is contained in some Λ_n. □

Note that the final construction of Λ given in theorem 25 was external to M, using the sequence $\langle \lambda_n \mid n < \omega \rangle$, although the graph was built as the union of initial segments that were elements of M and consequently Λ is at least amenable to M. But there seems little reason to expect for this construction that $\langle M, \in, \Lambda \rangle$ satisfies $\mathrm{ZFC}^-(\Lambda)$. Notice also that since λ is countable, the digraph Λ of theorem 25 is exactly the countable random λ-graded digraph. The important part of the theorem is that this digraph is set-like in M, which allows us to perform the modified Mostowski collapse of lemma 21.

Corollary 26. *If M is a countable transitive model of the theory $\mathrm{ZFC}^- + V = H_{\kappa^+}$ and $\kappa^{<\kappa} = \kappa$ in M and $\lambda = \mathrm{Ord}^M$, then there is a submodel $A \subseteq M$ such that $\langle A, \in \rangle$ is the countable random λ-graded digraph. Consequently, every countable λ-graded digraph is isomorphic to a submodel of $\langle A, \in \rangle$. In particular, every countable transitive model of set theory of height at most λ is isomorphic to a submodel of $\langle M, \in^M \rangle$.*

Proof. By theorem 25, there is a set-like λ-graded digraph Λ with the finite-pattern property, whose initial segments are sets in M. From this it follows that the modified Mostowski collapse of Λ as defined in lemma 21 is an isomorphism of $\langle \Lambda, \rightharpoonup \rangle \cong \langle A, \in \rangle$ for some $A \subseteq M$. Since Λ has the finite-pattern property, it and hence also $\langle A, \in \rangle$ are the countable random λ-graded digraph, which is universal for all countable λ-graded digraphs. Thus, every countable λ-graded digraph is isomorphic to a submodel of $\langle A, \in \rangle$ and hence of $\langle M, \in \rangle$. By observation 22, this includes all countable models of set theory of well-founded height at most λ. □

The key obstacle identified in the proof of observation 24 is the existence of more ordinals than the size of any given power set. The following theorem shows that this obstacle is the only obstacle, for if we are willing to give up the power set by adding many Cohen reals, we can find the desired universal digraph.

Theorem 27. *Assume* $V \models \mathrm{ZFC}$. *If* $G \subseteq \mathrm{Add}(\omega, \mathrm{Ord})$ *is* V*-generic for the forcing to add* Ord *many Cohen reals, then* $V[G] \models \mathrm{ZFC}^-$ *has a submodel* $A \subseteq V[G]$ *such that* $\langle A, \in \rangle$ *is an* Ord*-graded digraph with the finite-pattern property.*

Proof. The idea is to build a set-like Ord-graded digraph Λ inside $V[G]$. The nodes of Λ will be precisely the elements of $\mathrm{Ord} \times \omega$, and (α, n) will have value α. What remains is to specify the edges, which will be done generically. Let $\mathbb{P}$ be the class partial order consisting of a finite digraph G, whose vertices are contained in $\mathrm{Ord} \times \omega$, ordered by the (reverse) induced subgraph relation $G \leq H \leftrightarrow H \subseteq G$, so that as usual, stronger conditions are lower in the partial order.

I claim that this partial order is isomorphic to the forcing $\mathrm{Add}(\omega, \mathrm{Ord})$ to add Ord many Cohen reals, which is the same as the forcing to add a generic class function $g : \mathrm{Ord} \to 2$ by finite conditions. One can see this simply by observing that the forcing $\mathbb{P}$ is deciding generically and independently for each pair of nodes whether to place an edge $(\alpha, n) \rightharpoonup (\beta, k)$, where $\alpha < \beta$, or not. Thus, $\mathbb{P}$ is adding, by finite support, a generic function that decides of each possible edge, whether it is there or not. Thus, this forcing resonates with the random characterization of the countable random digraphs we saw earlier, but using genericity in place of randomness.

Suppose that Λ is the resulting V-generic Ord-grade digraph. It is clear that we will achieve the finite-pattern property, since if A, B and C are finite subsets of the nodes and α is an ordinal such that every node in A has value less than α and every node in B has value larger than α, then it is a dense requirement that there is a node with value α that is above all nodes in A, below all nodes in B and with no edges in C. The reason is that any condition mentions only finitely many additional nodes, and there will be an unmentioned node on level α that we may simply add in so as to satisfy the desired pattern.

Finally, the graph Λ is set-like in $V[G]$, because the predecessors of (α, n) in Λ are amongst $\alpha \times \omega$, which is a set. It follows that we may perform the modified Mostowski collapse of Λ as in lemma 21, and find $A \subseteq L[G]$ such that $\langle \Lambda, \rightharpoonup \rangle \cong \langle A, \in \rangle$. □

Corollary 28. *If* M *is a countable transitive model of set theory of height* λ*, then* M *is isomorphic to a submodel of* $\langle L_\lambda[G], \in \rangle$*, a forcing extension of* L_λ *obtained by adding* λ *many Cohen reals.*

Proof. By theorem 27, there is $A \subseteq L[G]$ such that $\langle A, \in \rangle$ is λ-graded and has the finite-pattern property. By theorem 14, it follows that $\langle A, \in \rangle$ is the countable random λ-graded digraph, which is universal for all countable λ-graded digraphs. Since $\langle M, \in\rangle$ is a λ-graded digraph by observation 22, it follows that $\langle M, \in\rangle$ is isomorphic to a submodel of $\langle A, \in\rangle$ and hence to a submodel of $\langle L[G], \in\rangle$. $\square$

This corollary will be improved by the main theorem, but it is already surprising, because the model M may have many large cardinals or other complex objects such as $0^\sharp$ that are fundamentally incompatible with $V = L$ or $V = L[G]$. Nevertheless, the main theorem will omit the need for G and obtain the even more surprising result that M actually embeds as a submodel of L_λ itself.

9. Proving the main theorems

Finally, I am ready to prove the main theorems 1 and 2.

Theorem 29 (Main Theorem 1)**.** *Every countable model of set theory $\langle M, \in^M\rangle$ is isomorphic to a submodel of its own constructible universe $\langle L^M, \in^M\rangle$. In other words, there is a quantifier-free-elementary embedding*

$$j : \langle M, \in^M\rangle \to \langle L^M, \in^M\rangle.$$

Proof. Suppose that M is a countable model of ZFC, not necessarily transitive or well-founded. Let $\lambda_n \in \mathrm{Ord}^M$ be an increasing cofinal sequence of limit ordinals from M. By the reflection theorem, we may assume without loss of generality that $V^M_{\lambda_n} \prec_{\Sigma_n} M$, although the argument will not really use much of this. Let Γ_n be the hypnagogic $(\lambda_n + 1)$-graded digraph as constructed in $L^M_{\lambda_{n+1}}$, running the construction for λ_{n+1} many birthdays. Thus, each value level of Γ_n has λ_{n+1} many elements, and Γ_n has a top layer consisting of the nodes with value λ_n. Furthermore, each Γ_n has the finite-pattern property, each Γ_n is an induced subgraph of Γ_{n+1} and the union $\Gamma = \bigcup_n \Gamma_n$ of the chain is exactly the hypnagogic Ord-graded digraph of L^M. Although each Γ_n is a set in L^M, the union digraph Γ is not set-like in L^M, since every node in Γ_n gains λ_{n+2} many new children in Γ_{n+1}.

I shall presently use the graphs Γ_n in order to construct a somewhat more elaborate graph, which I call the surrogate digraph, which will be set-like and which will exhibit a kind of finite-pattern property that will be sufficient for universality. Define that $\langle v_0, \ldots, v_n\rangle$ is a *surrogate sequence* of its final node v_n, if $v_n \in \Gamma_n$, the value of v_n is at least $\sup_{k<n} \lambda_k$ and each v_k for $k < n$ is a node in Γ_k of value λ_k. For surrogate sequences $v = \langle v_0, \ldots, v_n\rangle$ and

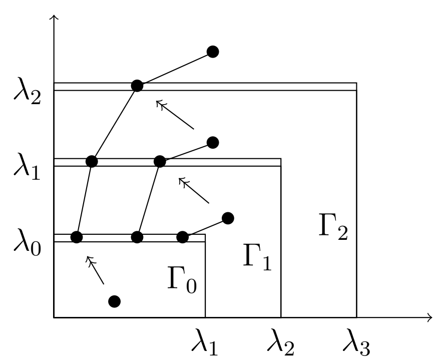

Figure 2. Surrogate sequences in the surrogate digraph

$w = \langle w_0, \ldots, w_m\rangle$, define the surrogate edge relation $w \twoheadrightarrow v$ to hold if and only if $m \le n$

and $w_m \rightharpoonup v_m$. The idea here is that we split the parent and child roles in Γ, so that only the final node w_m of a surrogate sequence $\langle w_0, \ldots, w_m \rangle$ acts as a child, while the prior nodes w_k act as surrogate parents, as far as gaining new children below value λ_k is concerned. It may be helpful to think of the relation that a set x has with its projections $x \cap V_\alpha$ for α of much smaller von Neumann rank than x, for we can tell if $y \in x$ or not, when y has rank less than α, merely by knowing whether or not $y \in x \cap V_\alpha$. In our final analysis, the surrogates v_k will be something like (but not exactly like) these projections, to capture the elements of rank less than λ_k, and to do so without leaving Γ_k. This surrogate maneuver will allow us to surmount the obstacle of observation 24.

Let Θ_n be the $(\lambda_n + 1)$-graded digraph consisting of all surrogate sequences of length at most n, using $\twoheadrightarrow$ as the edge relation, and giving the surrogate sequence the value of its final node in Γ. We will refer to the union digraph $\Theta = \bigcup_n \Theta_n$ as the Ord^M-graded *surrogate* digraph arising from $\langle \lambda_n \mid n < \omega \rangle$. Because we can define the sequence $\langle \lambda_n \mid n < \omega \rangle$ from Θ, we cannot expect that it is a class in M. Nevertheless, each Θ_n is a set in M and an induced subgraph of Θ, the restriction of Θ to surrogate sequences having value at most λ_n. It follows that the full surrogate digraph Θ is an amenable class over M and indeed over L^M. In particular, Θ is set-like in L^M, since any particular surrogate sequence $\vec{w}$ lies in some Θ_n, which is a set in L^M and is the $(\lambda_n + 1)$-initial segment of Θ.

Let's define that two subsets of Θ are *completely disjoint* if they are disjoint and furthermore any two surrogate sequences from either of them have no nodes in common. That is, not only are they disjoint as sets of sequences, but the individual nodes on those sequences do not recur.

Lemma 29.1. *The surrogate digraph Θ enjoys the surrogate finite-pattern property: if A, B and C are completely disjoint finite subsets of Θ and α is an ordinal of M, such that every element of A has value less than α and every element of B has value greater than α, then there is $v \in \Theta$ with value α such that $a \twoheadrightarrow v$ and $v \twoheadrightarrow b$ for all $a \in A$ and $b \in B$, and there are no $\twoheadrightarrow$ relations between v and elements of C, and furthermore, such that no nodes of v arise on any sequence from A, B or C.*

Proof. Let n be least such that $\alpha \leq \lambda_n$. We shall first choose the terminal node v_n of the desired surrogate sequence $v = \langle v_0, \ldots, v_n \rangle$. Let B_n consist of the nodes arising in the sequences of B with value in the interval $(\alpha, \lambda_n]$. These nodes arise either as the terminal nodes of sequences in B that terminate with a node of value at most λ_n, or else they arise as the surrogate nodes at level λ_n of a longer sequence. But in any case, $B_n \subseteq \Gamma_n$, and furthermore, for v to exhibit the correct $\twoheadrightarrow$ relation to the elements of B, it will suffice that $v_n \rightharpoonup b$ for each $b \in B_n$; and the main point of surrogates is that we shall be able to find such a v_n inside Γ_n, without care for the much larger value nodes that may appear later on in the sequences of B. But we must also ensure the correct $\twoheadrightarrow$ relation to A and C, so let A_n be the terminal nodes of any sequence in A having value at least $\sup_{k<n} \lambda_k$, and let C_n be the nodes occuring on a sequence in C that are in Γ_n. Since A, B and C are completely disjoint, it follows that A_n, B_n and C_n are disjoint subsets of Γ_n, and every node in A_n has value below α and every node in B_n has value above α. Thus, by the finite-pattern property of Γ_n, there is a new node $v_n \in \Gamma_n$ such that $a \rightharpoonup v_n$ and $v_n \rightharpoonup b$ for every $a \in A_n$ and $b \in B_n$ and such that v_n has no $\rightharpoonup$ relation with any node of C_n. Next, to define v_k for $k < n$, let A_k be the nodes arising from the sequences in A in the k^{th} layer, that is, with value in the interval $[\sup_{i<k} \lambda_i, \lambda_k)$, and similarly let C_k be any node on any sequence in C in Γ_k. (We

need not consider B for defining the surrogates v_k for $k < n$, since these surrogates act only as parents and not as children with respect to $\twoheadrightarrow$.) By the finite-pattern property of Γ_k, there is a new node v_k such that $a \rightharpoonup v_k$ for each $a \in A_k$ and v_k has no $\rightharpoonup$ relation with any node in C_k. This defines $v = \langle v_0, \ldots, v_n \rangle$, which has value α and which consists of entirely new nodes not arising in A, B and C. Furthermore, by design, $a \twoheadrightarrow v$ for each $a \in A$ and $v \twoheadrightarrow b$ for each $b \in B$, whilst v has no $\twoheadrightarrow$ relation with any sequence in C, and so we have fulfilled this instance of the surrogate finite-pattern property. □

So what we have is a set-like Ord^M-graded digraph $\Theta \subseteq L^M$, which does not have the finite-pattern property (and cannot by observation 6), but which has the surrogate finite-pattern property, and this is sufficient to carry out the universality construction. Specifically, we assign to each $x \in M$ a surrogate sequence v_x, with the same value as $\mathrm{rank}(x)$, ensuring that $x \in y \Leftrightarrow v_x \twoheadrightarrow v_y$. This can be achieved by enumerating $M = \{ x_n \mid n < \omega \}$, and choosing each v_{x_n} so as to relate via $\twoheadrightarrow$ to the previous v_{x_k} for $k < n$ in exactly the same way that x_n relates to x_k via $\in$, while also having the right value and inductively maintaining that the sequences v_{x_n} have no individual nodes in common. The surrogate finite-pattern property of the lemma exactly ensures that this recursive construction may proceed, and thus we build an embedding of $\langle M, \in^M \rangle$ into an induced subgraph of $\langle \Theta, \twoheadrightarrow \rangle$.

Finally, I argue that $\langle \Theta, \twoheadrightarrow \rangle \cong \langle A, \in^M \rangle$ for some $A \subseteq L^M$, using the fact that $\twoheadrightarrow$ is set-like in L^M. Specifically, since each Θ_n is a set in L^M, where $\twoheadrightarrow$ is well-founded, we may inside L^M carry out the modified Mostowski collapse of lemma 21, defining $\pi(v) = \{ \pi(w) \mid w \twoheadrightarrow v \} \cup \{ \{ \emptyset, v \} \}$ to find $\langle \Theta_n, \twoheadrightarrow \rangle \cong \langle A_n, \in^M \rangle \in L^M$. Furthermore, since each Θ_n is a $\twoheadrightarrow$ initial segment of Θ, it follows that these various isomorphisms cohere, and so $\langle \Theta, \twoheadrightarrow \rangle \cong \langle A, \in^M \rangle$, where $A = \bigcup_n A_n \subseteq L^M$.

Combining the two previous paragraphs, we observe that M is isomorphic to an induced subgraph of Θ, which is isomorphic to a submodel of L^M, and so M is isomorphic to a submodel of L^M, as desired. For the final claim of the theorem, note that an isomorphism of M to a submodel of L^M is precisely the same as a quantifier-free-elementary embedding $j : M \to L^M$. □

Corollary 30. *Every countable model $\langle M, \in^M \rangle$ of set theory is universal for all countable Ord^M-graded binary relations.*

Proof. This is what the proof of theorem 29 establishes. The surrogate finite-pattern property of Θ ensures that it is universal for all countable Ord^M-graded binary relations. □

Corollary 31. *A countable model $\langle M, \in^M \rangle$ of set theory embeds into another model $\langle N, \in^N \rangle$ of set theory if and only if the ordinals of M map order-preservingly into the ordinals of N.*

Proof. We needn't assume N is countable here, since we may simply pass to a countable submodel. For the forward implication, if $j : M \to N$ is an $\in$-embedding, then $f(\alpha) = \mathrm{rank}(j(\alpha))$ is an order-preserving map from Ord^M to Ord^N. Conversely, if the ordinals of M map into the ordinals of N, then we may externally view $\langle M, \in^M \rangle$ as an Ord^N-graded digraph, and the previous corollary shows that $\langle N, \in^N \rangle$ is universal for such relations. □

Corollary 32. *Every ill-founded countable model of set theory $\langle M, \in^M \rangle$ is universal for all countable acyclic binary relations.*

Proof. If M is not an ω-model, then already the hereditary finite sets HF^M of M are universal by theorem 9. If M has a standard ω, but the ordinals are ill-founded, then by results in

[Fri73], it follows that $\mathrm{Ord}^M \cong \lambda\cdot(1+\mathbb{Q})$ for some admissible ordinal λ. In particular, Ord^M contains a countable dense order, and so by lemma 4 every acyclic binary relation can be Ord^M-graded. Thus, $\langle M, \in^M\rangle$ is universal for all such relations. □

Theorem 33 (Main Theorem 2)**.** *The countable models of set theory are linearly pre-ordered by embeddability: for any two countable models of set theory $\langle M, \in^M\rangle$ and $\langle N, \in^N\rangle$, either M is isomorphic to a submodel of N or conversely. Indeed, the countable models of set theory are pre-well-ordered by embeddability in order type exactly $\omega_1 + 1$.*

Proof. Corollary 32 shows that the models having ill-founded ordinals are all universal, so that every countable model of set theory embeds into any of them. In particular, they are all bi-embeddable with each other. What remains are the well-founded models, which by corollary 31 are ordered the same as their ordinals. In particular, any two well-founded models with the same ordinals are bi-embeddable, and any shorter model is embeddable into any taller model.

Thus, the bi-embeddability classes of the models of set theory are in one-to-one correspondence with the collection of ordinals λ such that L_λ is a model of set theory, plus one more bi-embeddability class at the top for the ill-founded models. Since there are ω_1 many countable admissible ordinals, it follows that the order-type of the bi-embeddability classes of the countable models of set theory (understood as models of KP) is precisely $\omega_1 + 1$. □

If one wants to consider only the countable models of ZFC, or only the models of ZFC plus large cardinals, then the argument shows that they are pre-well-ordered in order type $\eta+1$, where η is the number of countable ordinals λ for which there is a model of the theory having height λ. If there is an inaccessible cardinal, for example, then it follows that $\eta = \omega_1$ for ZFC models and the countable models of ZFC are pre-well-ordered by embeddability in order type $\omega_1 + 1$.

Finally, let me observe that one definitely doesn't have linearity for the embeddability relation in the context of uncountable models.

Observation 34. *If* ZFC *is consistent, then there are two models $M, N \models$ ZFC, such that neither is isomorphic to a submodel of the other.*

Proof. If ZFC is consistent, then on the one hand, there is an uncountable model $M \models$ ZFC with cofinality ω and with every cut of cofinality ω. On the other hand, one may also build an ω_1-like model $N \models$ ZFC, so N is uncountable, but every initial segment is countable. It follows that N has cofinality ω_1.

Note that M cannot embed into N, since there are elements of M with uncountably many predecessors, but every element of N has only countably many predecessors. And N cannot embed into M, since the supremum of the image would be a cut of uncountable cofinality. □

10. Questions

Although the main theorem shows that every countable model of set theory embeds into its own constructible universe

$$j : M \to L^M,$$

this embedding j is constructed completely externally to M and there is little reason to expect that j could be a class in M or otherwise amenable to M. To what extent can we prove or refute the possibility that j is a class in M? This amounts to considering the matter internally as a question about V. Surely it would seem strange to have a class embedding $j : V \to L$ when $V \neq L$, even if it is elementary only for quantifier-free assertions, since such an embedding is totally unlike the sorts of embeddings that one usually encounters in set theory. Nevertheless, I am at a loss to refute the hypothesis, and the possibility that there might be such an embedding is intriguing, if not tantalizing, for one imagines all kinds of constructions that pull structure from L back into V.

Question 35. *Can there be an embedding $j : V \to L$ when $V \neq L$?*

By embedding, I mean an isomorphism from $\langle V, \in\rangle$ to its range in $\langle L, \in\rangle$, which is the same as a quantifier-free-elementary map $j : V \to L$, a map for which $x \in y \Leftrightarrow j(x) \in j(y)$. (Note that if $V = L$, then we have some easily defined non-identity embeddings $j : L \to L$ arising from lemma 21, such as $j(x) = \{\, j(y) \mid y \in x \,\} \cup \{\, \{\, \emptyset, x \,\} \,\}$, which has $j(x) \neq x$ for every x.) Question 35 is most naturally formalized in Gödel-Bernays set theory, asking whether there can be a GB-class j forming such an embedding. If one wants $j : V \to L$ to be a definable class, then this of course implies $V = \mathrm{HOD}$, since the definable L-order can be pulled back to V, via $x \leq y \Leftrightarrow j(s) \leq_L j(y)$. More generally, if j is merely a class in Gödel-Bernays set theory, then the existence of an embedding $j : V \to L$ implies global choice, since from the class j we can pull back the L-order. For these reasons, we cannot expect every model of ZFC or of GB to have such embeddings. Can they be added generically? Do they have some large cardinal strength? Are they outright refutable?

It they are not outright refutable, then it would seem natural that these questions might involve large cardinals; perhaps $0^\sharp$ is relevant. But I am unsure which way the answers will go. The existence of large cardinals provides extra strength, but may at the same time make it harder to have the embedding, since it pushes V further away from L. For example, it is conceivable that the existence of $0^\sharp$ will enable one to construct the embedding, using the Silver indiscernibles to find a universal submodel of L; but it is also conceivable that the non-existence of $0^\sharp$, because of covering and the corresponding essential closeness of V to L, may make it easier for such a j to exist. Or perhaps it is simply refutable in any case. The first-order analogue of the question is:

Question 36. *Does every set A admit an embedding $j : \langle A, \in\rangle \to \langle L, \in\rangle$? If not, which sets do admit such embeddings?*

The main theorem shows that every countable set A embeds into L. What about uncountable sets? Let us make the question extremely concrete:

Question 37. *Does $\langle V_{\omega+1}, \in\rangle$ embed into $\langle L, \in\rangle$? How about $\langle P(\omega), \in\rangle$ or $\langle \mathrm{HC}, \in\rangle$?*

It is also natural to inquire about the nature of $j : M \to L^M$ even when it is not a class in M. For example, can one find such an embedding for which $j(\alpha)$ is an ordinal whenever α is an ordinal? The embedding arising in the proof of theorem 29 definitely does not have this feature, since every set in M is mapped to a node in Θ, which is mapped ultimately to $\pi(v) = \{\, \pi(w) \mid w \twoheadrightarrow v \,\} \cup \{\{\emptyset, v\}\}$, and this latter set is never an ordinal.

Question 38. *Does every countable model $\langle M, \in^M\rangle$ of set theory admit an embedding $j : M \to L^M$ that takes ordinals to ordinals?*

Probably one can arrange this simply by being a bit more careful with the modified Mostowski procedure in lemma 21, as it is applied to Θ in the proof of the main theorem,

theorem 29. And if this is correct, then numerous further questions immediately come to mind, concerning the extent to which we ensure more attractive features for the embeddings j that arise in the main theorems. This will be particularly interesting in the case of well-founded models, as well as in the case of $j : V \to L$, as in question 35, if that should be possible.

Question 39. *Can there be a nontrivial embedding $j : V \to L$ that takes ordinals to ordinals?*

Finally, we inquire about the extent to which the main theorems of this article can be extended from the countable models of set theory to the ω_1-like models:

Question 40. *Does every ω_1-like model of set theory $\langle M, \in^M \rangle$ admit an embedding $j : M \to L^M$ into its own constructible universe? Are the ω_1-like models of set theory linearly pre-ordered by embeddability?*

References

[Acz88] Peter Aczel. *Non-well-founded sets*. Center for the Study of Language and Information, Stanford, CA, 1988.

[Del11] Ewan Delanoy. Comparing countable models of ZFC. Mathematics, 2011. URL:http://math.stackexchange.com/q/64095 (version: 2011-09-13).

[ESV11] A. Enayat, J. Schmerl, and A. Visser. ω-models of finite set theory. In J. Kennedy and R. Kossak, editors, *Set theory, Arithmetic, and Foundations of Mathematics: Theorems, Philosophies*, number 36 in Lecture Notes in Logic, chapter 4. Cambridge University Press, 2011.

[Fri73] Harvey Friedman. Countable models of set theories. *Cambridge Summer School in Mathematical Logic (1971), Lecture Notes in Math.*, 337:pp. 539–573, 1973.

[Hau02] Felix Hausdorff. *Felix Hausdorff—gesammelte Werke. Band II.* Springer-Verlag, Berlin, 2002. “Grundzüge der Mengenlehre”. [“Foundations of set theory”], Edited and with commentary by E. Brieskorn, S. D. Chatterji, M. Epple, U. Felgner, H. Herrlich, M. Hušek, V. Kanovei, P. Koepke, G. Preuß, W. Purkert and E. Scholz.

[Hod93] Wilfrid Hodges. *Model theory*, volume 42 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1993.

[Kir10] Laurence Kirby. Substandard models of finite set theory. *Mathematical Logic Quarterly*, 56(6):631–642, 2010.

[KW07] Richard Kaye and Tin Lok Wong. On interpretations of arithmetic and set theory. *Notre Dame J. Formal Logic*, 48(4):497–510, 2007.

[McA82] Kenneth McAloon. On the complexity of models of arithmetic. *The Journal of Symbolic Logic*, 47(2):pp. 403–415, 1982.

[Par84] Jeff B. Paris. On the structure of models of bounded E_1-induction. *Časopis Pěst. Mat.*, 109(4):372–379, 1984.

[Res86] J. P. Ressayre. Introduction aux modèles récursivement saturés. In *Séminaire Général de Logique 1983–1984 (Paris, 1983–1984)*, volume 27 of *Publ. Math. Univ. Paris VII*, pages 53–72. Univ. Paris VII, Paris, 1986.

Mathematics Program, The Graduate Center of The City University of New York, 365 Fifth Avenue, New York, NY 10016, U.S.A. & Department of Mathematics, The College of Staten Island of CUNY, Staten Island, NY 10314, U.S.A.

E-mail address: jhamkins@gc.cuny.edu

URL: http://jdh.hamkins.org