

The metamathematics of ergodic theory

Jeremy Avigad

*Department of Philosophy and Department of Mathematical Sciences
Carnegie Mellon University
Pittsburgh, PA 15213*

Abstract

The metamathematical tradition, tracing back to Hilbert, employs syntactic modeling to study the methods of contemporary mathematics. A central goal has been, in particular, to explore the extent to which infinitary methods can be understood in computational or otherwise explicit terms. Ergodic theory provides rich opportunities for such analysis. Although the field has its origins in seventeenth century dynamics and nineteenth century statistical mechanics, it employs infinitary, non-constructive, and structural methods that are characteristically modern. At the same time, computational concerns and recent applications to combinatorics and number theory force us to reconsider the constructive character of the theory and its methods. This paper surveys some recent contributions to the metamathematical study of ergodic theory, focusing on the mean and pointwise ergodic theorems and the Furstenberg structure theorem for measure preserving systems. In particular, I characterize the extent to which these theorems are nonconstructive, and explain how proof-theoretic methods can be used to locate their “constructive content.”

Key words: ergodic theory, computability, constructive mathematics

1 Introduction

The late nineteenth century inaugurated an era of sweeping changes in mathematics. Whereas mathematics had, until that point, been firmly rooted in explicit construction and symbolic calculation, the new developments emphasized a kind of understanding that was often at odds with computational

Email address: avigad@cmu.edu (Jeremy Avigad).

URL: <http://www.andrew.cmu.edu/user/avigad> (Jeremy Avigad).

¹ Work partially supported by NSF grant DMS-0700174.

concerns. Dedekind, for example, wrote of his development of the theory of ideals:

It is preferable, as in the modern theory of functions, to seek proofs based immediately on fundamental characteristics, rather than on calculation, and indeed to construct the theory in such a way that it is able to predict the results of calculation. . . [21, page 102]

Such attitudes paved the way to the adoption of the infinitary, nonconstructive, set theoretic, algebraic, and structural methods that are characteristic of modern mathematics.

The new methods were controversial, however. At issue was not just whether they are consistent, but, more pointedly, whether they are meaningful, and appropriate to mathematics. After all, if one views mathematics as an essentially computational science, then arguments without computational content, whatever their heuristic value, are not properly mathematical. The discovery of logical and set-theoretic paradoxes at the turn of the twentieth century, however, brought the issue of consistency to the fore, and Brouwer's intuitionistic challenges in the 1910's made the problem of finding an adequate foundation and justification of the new methods even more pressing.

David Hilbert's metamathematical program, unveiled in 1922 [42,43], is often viewed, narrowly, as a response to the consistency problem, but it is better seen as an attempt to justify the new methods in light of the broader concerns just described. The strategy Hilbert proposed was to model the new methods using formal, axiomatic systems, and then prove the consistency of those systems using "finitary" methods, whose validity could not be questioned. Within a formal system, one can enjoy modern methods to one's heart's content, with the knowledge that from a metamathematical standpoint, the symbolic rules endow the resulting proofs with an explicit combinatorial content. Of course, something more is needed to justify the choice of symbolic rules with respect to our understanding of the mathematical enterprise; at the bare minimum, we wish to know that the universal assertions we derive in the system will not be contradicted by our experiences, and the existential predictions will be born out by calculation. This is exactly what Hilbert's program was designed to do.

Kurt Gödel's incompleteness theorems of 1931 [36] demonstrated the impossibility of achieving Hilbert's goal, assuming that the safe, finitary portion of mathematics is included among the broader range of methods to be justified. But the more general program of understanding modern methods in syntactic terms, and using that understanding to clarify their computational content, has been more successful. For one thing, we now know that significant portions of mathematics can be formalized in theories that are strictly

weaker than primitive recursive arithmetic, and therefore have a finitary justification.² Proof-theoretic methods now provide numerous ways of “reducing” classical theories to constructive ones [3,24,25], and “mining” the constructive content of classical proofs [34,51].

The theory of dynamical systems and ergodic theory provide fruitful arenas for such analysis. Although these subjects arose from the study of physical and statistical phenomena, they make full use of modern structural methods that do not directly bear on the original computational concerns. My goal here will be to survey some recent developments in the metamathematics of ergodic theory with these issues in mind. In particular, I will try to clarify the extent to which the methods of ergodic theory can be given a direct computational interpretation, and explain how proof-theoretic methods enable us to obtain useful information in situations when the methods are explicitly nonconstructive. Most of the work I will describe here has been carried out jointly with Philipp Gerhardy, Ksenija Simic, and Henry Towsner.

2 Dynamical systems and ergodic theory

A *discrete dynamical system* consists of a structure, \mathcal{X} , and a map T from (the underlying set of) \mathcal{X} to itself. One can view \mathcal{X} as a space of configurations, or states, of a physical system that evolves over time. Assuming x is any such state, one can take Tx to be the state of the system after one unit of time has elapsed. To have anything interesting to say, one has to assume \mathcal{X} bears some structure; for example, \mathcal{X} may be a metric space, a topological space, or a differentiable manifold. Laplace took this general model to form the scientific basis for his mechanistic view of the universe:

We ought then to regard the present state of the universe as the effect of its previous state and as the cause of that which is to follow. An intelligence that, at a given instant, could comprehend all the forces by which nature is animated and the respective situation of all the things that make it up, if moreover it were vast enough to submit these data to analysis, would encompass in the same formula the movements of the greatest bodies of the universe and those of the lightest atoms. For such an intelligence, nothing would be uncertain, and the future, as the past, would be present to its eyes [59, page 2].

The study of dynamical systems can be traced back to Newton, who essentially solved the two-body problem in providing a closed-form determination

² See Simpson [66] and Avigad [4] for more precise claims, and Burgess [18] for a helpful caveat.

of the behavior of two bodies whose motion is constrained only by the gravitational force between them. Newton soon learned that when one adds a third body, the situation becomes significantly more difficult; he reported to the astronomer John Machin that “his head never ached but with his studies on the moon” (quoted in [76, page 544], and in [9, page 15]). Some of the greatest mathematical minds of the eighteenth and nineteenth century, including Euler, Lagrange, and Jacobi, were caught up in the three-body problem, and the heady optimism of the Laplacian world view was soon thwarted. It was Poincaré who first caught a glimpse of the modern theory of chaos, with the realization that part of the problem lies in the sensitivity of a system to its initial conditions:

If we knew exactly the laws of nature and the situation of the universe at the initial moment, we could predict exactly the situation of that same universe at a succeeding moment. But even if it were the case that the natural laws had no longer any secret for us, we could still only know the initial situation *approximately*. If that enabled us to predict the succeeding situation with *the same approximation*, that is all we require, and we should say that the phenomenon had been predicted, that it is governed by laws. But it is not always so; it may happen that small differences in the initial conditions produce very great ones in the final phenomena. A small error in the former will produce an enormous error in the latter. Prediction becomes impossible, and we have the fortuitous phenomenon [62, page 68].

The beginnings of a breakthrough came in the 1880’s when Poincaré shifted attention from the “quantitative” features of curves defined by differential equations to more “qualitative” features.

... this qualitative study has in itself an interest of the first order. Several very important questions of analysis and mechanics reduce to it. Take for example the three body problem: one can ask if one of the bodies will always remain within a certain region of the sky or even if it will move away indefinitely; if the distance between two bodies will infinitely decrease or diminish, or even if it will remain within certain limits. Could one not ask a thousand questions of this type which would be resolved when one can construct qualitatively the trajectories of the three bodies? ([61]; translation [9, page 31])

With this work, the focus of the theory of dynamical systems moved to ways of characterizing their global behavior, even in situations that are not amenable to closed-form solution or computational approximation.

This essay will focus on *ergodic theory*, where the system \mathcal{X} is assumed to be a finite measure space (X, \mathcal{B}, μ) such that $\mu(X)$ is finite and T is assumed to be a measure preserving transformation, which is to say, T satisfies $\mu(T^{-1}A) =$

$\mu(A)$ for every $A \in \mathcal{B}$. We will be solely concerned with cases where \mathcal{X} is separable, and I will refer to such a system as a *measure preserving system*. These are often used to model physical processes; for example, if we model the state of a system of N particles by giving the position and momentum of each particle, the evolution of the system, in accordance with Hamilton's equations, preserves Lebesgue measure. Measure preserving systems can also be used to model probabilistic processes: for each $A \in \mathcal{B}$, $\mu(A)$ is the probability that the system is in a state in A at time 0, in which case $\mu(T^{-1}A)$ is the probability that the system will be in a state in A one unit of time later.³ We will see, below, that measure preserving systems have useful applications in number theory and combinatorics, where the measures in question are carefully tailored to the application.

Although ergodic theory has its roots in seventeenth century dynamics and nineteenth century statistical mechanics, the field is quintessentially modern, enjoying the full range of algebraic, infinitary, nonconstructive, and structural methods. These concerns are sometimes at odds with the motivating computational concerns, a tension we will explore below.

3 Analysis of the ergodic theorems

Let $\mathcal{X} = (X, \mathcal{B}, \mu, T)$ be a measure preserving system, and let $f \in L^1(\mathcal{X})$ be any real-valued integrable function. If we think of f as representing the result of performing a measurement on the state of the system, then $f \circ T$ represents the result of performing that same measurement after one unit of time. Starting in state x , suppose now we perform n measurements,

$$f(x), f(Tx), f(T^2x), \dots, f(T^{n-1}x),$$

and take their average. In the long run, do these averages stabilize?

Note that T induces an isometry $\hat{T} : f \mapsto f \circ T$ of the space $L^1(\mathcal{X})$ of integrable functions from \mathcal{X} to \mathbb{R} , and of the Hilbert space $L^2(\mathcal{X})$ of square-integrable functions. For each $n \geq 1$, define $A_n f$ to be the function $\frac{1}{n} \sum_{i < n} \hat{T}^i f$, so that, for each x , $A_n f(x)$ denotes on the average measurement over the first n points in the orbit of x . The von Neumann mean ergodic theorem [73] asserts that for any f in $L^2(\mathcal{X})$, the sequence $(A_n f)$ converges in the L^2 norm. The Birkhoff pointwise ergodic theorem [12] asserts that, moreover, for any f in $L^1(\mathcal{X})$, the sequence $(A_n f)$ converges pointwise, almost everywhere, and in the L^1

³ Indeed, it is often not acknowledged that this is the implicit context of the Laplace quotation above. According to Laplace, since the evolution of a dynamical system is completely determined by its state, it is merely our ignorance of the precise state that forces us to resort to probabilistic notions.

norm. A clean geometric proof due to Riesz [64] shows that the von Neumann theorem holds more generally for any nonexpansive operator \hat{T} on a Hilbert space, that is, any operator satisfying $\|Tf\| \leq \|f\|$ for every f .

The measure preserving system \mathcal{X} is said to be *ergodic* if it cannot be decomposed into nontrivial components that are invariant under T , that is, if $T^{-1}A = A$ implies that $\mu(A) = 0$ or $\mu(A) = 1$. When \mathcal{X} is ergodic, the mean and pointwise ergodic theorems imply that $(A_n f)$ converge to the constant function $\int f d\mu$, in all the senses of convergence indicated above. In other words, if the space is ergodic, the result of averaging a measurement over time is, in the limit, equivalent to averaging the measurement over all possible configurations of the system.

It is now reasonable to ask how quickly the sequence $(A_n f)$ converges, and whether a bound on the rate of convergence can be computed from the initial data. In other words, given T and f , can one compute a function $r : \mathbb{Q} \rightarrow \mathbb{N}$ such that for every rational $\varepsilon > 0$, $\|A_m f - A_{r(\varepsilon)} f\| < \varepsilon$ whenever $m \geq r(\varepsilon)$? It is known that a sequence of ergodic averages can converge arbitrarily slowly (see [49,58] for precise formulations of this statement and related results), but the question as to the computability of an r from the initial data is a separate issue. For example, if $(a_n)_{n \in \mathbb{N}}$ is any sequence of real numbers that decreases to 0, no matter how slowly, one can compute a bound on the rate of convergence by systematically querying the elements of the sequence until one of them is seen to drop below ε . On the other hand, it is not hard to construct a computable sequence $(b_n)_{n \in \mathbb{N}}$ of rational numbers that converges to 0, with the property that no computable function $r(\varepsilon)$ meets the specification above. In a similar way, one can construct a computable sequence $(c_n)_{n \in \mathbb{N}}$ of rational numbers that is monotone and bounded, but converges to a noncomputable real number. Thus, neither monotonicity nor the existence of a computable limit alone is enough to guarantee the effective convergence of a sequence of rationals.

What these examples show is that the question as to whether it is possible to compute a bound on a rate of convergence of a sequence from some initial data is not a question about the speed of the sequence's convergence, but, rather, its predictability. To make the question precise, one needs to rely on standard notions of computability in analysis; see [6,63,75] for details. Simic and I [7,65] have shown that, in general, one cannot compute a bound on the rate of convergence from the initial data; the following formulation is taken from [6].⁴

Theorem 3.1 *There are a computable measure preserving transformation of*

⁴ Jan Reimann has recently brought to my attention work by V. V. V'yugin [74], which also establishes the noncomputability of rates of convergence in the ergodic theorems.

$[0, 1]$ under Lebesgue measure and a computable characteristic function $f = \chi_A$ such that if $f^* = \lim_n A_n f$, then $\|f^*\|_2$ is not a computable real number. In particular, f^* is not a computable element of $L^2(\mathcal{X})$, and there is no computable bound on the rate of convergence of $(A_n f)$ in either the L^2 or L^1 norm.

On the other hand, Gerhardy, Towsner, and I [6] have shown that to compute a bound on the rate of convergence, and hence the limit of the sequence, it suffices to know the norm of the limit.

Theorem 3.2 *Let \hat{T} be a nonexpansive operator on a separable Hilbert space and let f be an element of that space. Let $f^* = \lim_n A_n f$. Then f^* , and a bound on the rate of convergence of $(A_n f)$ in the Hilbert space norm, can be computed from f , \hat{T} , and $\|f^*\|$. In particular, if \hat{T} arises from an ergodic transformation T , then f^* is computable from T and f .*

The second statement follows from the first, since in any ergodic space the averages $(A_n f)$ converge to the constant function equal to $\int f d\mu$, which is computable from f .

The negative result of Theorem 3.1 is not surprising. The ergodic theorem deals with the limiting behavior of dynamical systems, and one would not expect such limiting behavior to be computable in every case. After all, the “limiting behavior” of a Turing machine should include a determination as to whether or not the machine halts on a given input, which is the most basic example of an undecidable problem. From a logical perspective, the assertion that the sequence $(A_n f)$ converges can be represented as follows:

$$\forall \varepsilon > 0 \exists n \forall m > n (\|A_m f - A_n f\| < \varepsilon). \quad (1)$$

It is the inner universal quantifier that makes it impossible to compute a witness to the existential quantifier, since, in general, there is no finite test one can perform to determine whether a given n has the requisite property. But although Theorem 3.1 could have been anticipated, it is somewhat disconcerting. What good is a convergence theorem if, in general, we cannot determine the rate of convergence? Is there any constructive information to be had?

Bishop [13–15] provides one answer. The assertion that a bounded sequence (a_n) converges is classically equivalent to the assertion that for every $\alpha < \beta$, the sequence crosses the strip between α and β at most finitely many times. (To see this, note that $\liminf a_n$ and $\limsup a_n$ are always defined, with $\liminf a_n \leq \limsup a_n$; the condition rules out a strict inequality.) Bishop used this idea to fashion a constructive version of the pointwise ergodic theorem, which implies its classical counterpart. There has lately been a resurgence of interest in such “upcrossing inequalities”; see [44,45,47–50].

Gerhardy, Towsner, and I [6] provide an alternative approach. Assertion (1)

is classically equivalent to the assertion that for any function K from \mathbb{N} to \mathbb{N} , the following holds:

$$\forall \varepsilon > 0 \exists n \forall m \in [n, K(n)] \|A_m f - A_n f\| < \varepsilon. \quad (2)$$

Given $\varepsilon > 0$, clearly any n witnessing (1) satisfies (2) for any K . Conversely, if (1) were false, then for some $\varepsilon > 0$ and every n , one could find an $m > n$ such that $\|A_m f - A_n f\| \geq \varepsilon$. Letting K be the function that for each $n \geq 1$ returns such an m yields a counterexample to (2). This yields our constructive version of the mean ergodic theorem:

Theorem 3.3 *Let T be any nonexpansive mapping on a Hilbert space, let f be any element of that space, let $\varepsilon > 0$, and let K be any function. Then there is an $n \geq 1$ such that for every m in $[n, K(n)]$, $\|A_m f - A_n f\| < \varepsilon$.*

A special case of this statement has recently been used by Tao [70]. Gerhardy, Towsner, and I provide a constructive proof of Theorem 3.3, with explicit bounds on n expressed solely in terms of K and $\rho = \|f\|/\varepsilon$. In particular, our bounds are uniform on any ball in Hilbert space and independent of T . As special cases, we have the following:

- If $K = n^{O(1)}$, then $n(f, \varepsilon) = 2^{2^{O(\rho^2 \log \log \rho)}}$.
- If $K = 2^{O(n)}$, then $n(f, \varepsilon) = 2_{O(\rho^2)}^1$.
- If $K = O(n)$ and T is an isometry, then $n(f, \varepsilon) = 2^{O(\rho^2 \log \rho)}$.

Similar considerations hold for the pointwise ergodic theorem, which is classically equivalent to the following:

Theorem 3.4 *Given T and f as above, for every $\lambda_1 > 0$, $\lambda_2 > 0$, and K , there is an $n \geq 1$ satisfying*

$$\mu(\{x \mid \max_{n \leq m \leq K(n)} |A_n f(x) - A_m f(x)| > \lambda_1\}) \leq \lambda_2.$$

For f in $L^2(\mathcal{X})$, Gerhardy, Towsner, and I provide explicit bounds on n in terms of f , λ_1 , λ_2 , and K .

Our noncomputability result, Theorem 3.1, can be relativized, and yields precise information as to the degrees of noncomputability of ergodic limits.⁵ The results can also be cast in terms of provability in weak or constructive axiomatic frameworks. Details are spelled out in [6], and stronger results on the

⁵ Specifically, if the measure preserving system and the function f are computable from a set A , then the limit of the averages $(A_n f)$ is computable from A' ; and for every A , there are a system and function computable from A such that A' is computable from the limit of the averages.

reverse mathematics of the ergodic theorems can be found in [7,65].⁶ It is worth noting that our constructive version of the mean ergodic theorem is an example of Kreisel’s no-counterexample interpretation [55,56], and our extractions of bounds can be viewed as applications of a body of proof theoretic results that fall under the heading “proof mining” (see, for example, [34,51]).

4 The Furstenberg structure theorem and ergodic Ramsey theory

If a measure preserving system $\mathcal{X} = (X, \mathcal{B}, \mu, T)$ is ergodic, the pointwise ergodic theorem implies that the space has a certain “mixing” property: almost every orbit x, Tx, T^2x, \dots traverses the space with enough regularity so that for every integrable function f , the average of f over the sequence is equal to the average measurement over the entire space.⁷ Ergodicity is also equivalent to saying that for every pair of measurable sets A and B , we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mu(T^{-i}A \cap B) = \mu(A)\mu(B).$$

This says, roughly, that the probability of being in a state in B at time 0 and in a state in A after i units of time, is, on the average, close to the product of the probability of being in A and the probability of being in B , assuming the average is taken over a sufficiently large period of time.

Ergodicity is not a very strong mixing property, and any space can be decomposed into ergodic components. A space is said to be (*strong*) *mixing* if for every A and B , one has

$$\lim_{n \rightarrow \infty} \mu(T^{-n}A \cap B) = \mu(A)\mu(B).$$

This means that for sufficiently large n the probability of being in B at time 0 and then in A at time n is roughly the product of the individual probabilities.

⁶ For example, the Riesz proof of the mean ergodic theorem shows that if T is any nonexpansive map on a Hilbert space \mathcal{H} , then \mathcal{H} can be decomposed as an orthogonal sum of the subspace $\mathcal{M} = \{f \mid Tf = f\}$ of fixed points, and the subspace \mathcal{N} that is the closure of the set spanned by vectors of the form $\{Tf - f\}$. It is then easy to show that the ergodic averages $A_n f$ converge to the projection of f on \mathcal{M} . Simic and I [7] show that, over RCA_0 , the statement that $(A_n f)$ converges is equivalent to the assertion that the projection of f on \mathcal{N} exists; but, surprisingly, the statement “if the projection of f on \mathcal{M} exists then $A_n f$ converges” is still equivalent to arithmetic comprehension.

⁷ The pointwise ergodic theorem implies that this is true for a single function f ; for the stronger claim just made, note that the conclusion can be made to hold for a countable dense set of functions, simultaneously.

So being mixing means being random in the sense that over sufficiently long periods of time, events are uncorrelated: knowing that we are in a state in B at time 0 does not give much information about what states we might be in at a later time n . Although it is not readily apparent, a more natural and better-behaved property of a system is that of being *weak mixing*:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n |\mu(T^{-i}A \cap B) - \mu(A)\mu(B)| = 0.$$

This turns out to be equivalent to saying that $\lim_{n \rightarrow \infty} \mu(T^{-n}A \cap B) = \mu(A)\mu(B)$ holds for every A and B , once we exclude a set of natural numbers of 0 density.⁸

A system is said to be *compact* if it has the property that for every f in $L^2(X, \mathcal{B}, \mu)$, the orbit $\{f, \hat{T}f, \hat{T}^2f, \dots\}$ has compact closure. This is equivalent to saying that for every such f , the set $\{f, \hat{T}f, \hat{T}^2f, \dots\}$ is *totally bounded*: for every $\varepsilon > 0$, there is an n , such that every $\hat{T}^i f$ is within a distance of ε of $\{f, \hat{T}f, \dots, \hat{T}^n f\}$ in the L^2 norm.

Taking f to be the characteristic function of a set, compactness implies, roughly, that events tend to recur at regular intervals. Thus compactness and weak mixingness characterize opposite behaviors: a compact system exhibits a high degree of regularity and order, while a weak mixing system exhibits a high degree of randomness. This opposition is fundamental in analysis: a system can be rigid, or chaotic; a channel can carry signal, or noise. In general, a system will be neither weak mixing nor compact. However, a remarkable theorem, due to Furstenberg [29,30,32], provides a structural decomposition of any system in terms of these two types of behavior. First, we present two key lemmas.

Lemma 4.1 (Koopman and von Neumann [52]) *If a measure preserving system is not weak mixing, it has a nontrivial compact T -invariant factor.*

There are a number of equivalent ways to think of a factor. If (X, \mathcal{B}, μ) is a measure space, one way to present a factor of the system is simply to provide a sub- σ -algebra of sets $\mathcal{B}' \subseteq \mathcal{B}$. Thus, the factor (X, \mathcal{B}', μ) is a coarsening of the original system that can “see” fewer events. At the extreme, the trivial factor consists of the two events $\{\emptyset, X\}$. Any homomorphism f from a space (X, \mathcal{B}, μ) to a space (Y, \mathcal{C}, ν) gives rise to the factor (X, \mathcal{B}', μ) with $\mathcal{B}' = f^{-1}\mathcal{C}$, and, conversely, every factor is of this form; thus one can also view a factor as a quotient, or homomorphic image, of the initial space. A factor $\mathcal{B}' \subseteq \mathcal{B}$ also gives rise to the closed subspace $L^2(X, \mathcal{B}', \mu)$ of $L^2(X, \mathcal{B}, \mu)$ that contains the constants and is closed under min and max, and, once again, every factor arises

⁸ It is also equivalent to saying that the product, $(X \times X, \mathcal{B} \times \mathcal{B}, \mu \times \mu, T \times T)$ is ergodic.

in this way. By definition, the factor is T -invariant when the sub- σ -algebra is invariant under T ; equivalently, when the homomorphism is a homomorphism of measure-preserving systems, or when the closed subspace L^2 is \hat{T} -invariant.

One can define “relativized” notions of compactness and weak mixing in such a way that the following generalization of Lemma 4.1 holds:

Lemma 4.2 (Furstenberg [29]) *If a measure preserving system (X, \mathcal{B}, μ, T) is not weak mixing relative to a proper T -invariant factor \mathcal{B}_1 , there is a T -invariant factor \mathcal{B}_2 , such that $\mathcal{B}_1 \subsetneq \mathcal{B}_2 \subseteq \mathcal{B}$ and $(X, \mathcal{B}_2, \mu, T)$ is compact relative to $(X, \mathcal{B}_1, \mu, T)$.*

Of course, if $\mathcal{B}_2 \neq \mathcal{B}$ and (X, \mathcal{B}, μ, T) is not weak mixing relative to \mathcal{B}_2 , we can repeat the process and find another intermediate factor $\mathcal{B}_2 \subsetneq \mathcal{B}_3 \subseteq \mathcal{B}$. This can be iterated as long as the hypothesis of the lemma holds, yielding a sequence of factors $\mathcal{B}_1 \subsetneq \mathcal{B}_2 \subsetneq \dots \subseteq \mathcal{B}$. We can continue the process into the transfinite by taking unions at limit stages $\mathcal{B}_\lambda = \cup_{\alpha < \lambda} \mathcal{B}_\alpha$. This gives rise to a strictly increasing sequence of subspaces $L^2(X, \mathcal{B}_\alpha, \mu)$ of a separable Hilbert space, $L^2(X, \mathcal{B}, \mu)$, and so the process has to stop at some countable ordinal α .⁹ Thus, we have the following:

Theorem 4.3 (Furstenberg structure theorem) *Let (X, \mathcal{B}, μ, T) be any ergodic measure preserving system. Then there is a transfinite increasing sequence of factors $(\mathcal{B}_\alpha)_{\alpha \leq \gamma}$ such that:*

- (1) \mathcal{B}_0 is the trivial factor, $\{\emptyset, X\}$
- (2) For each $\alpha < \gamma$, $(X, \mathcal{B}_{\alpha+1}, \mu, T)$ is compact relative to $(X, \mathcal{B}_\alpha, \mu, T)$.
- (3) For each limit $\lambda \leq \gamma$, $\mathcal{B}_\lambda = \cup_{\alpha < \lambda} \mathcal{B}_\alpha$.
- (4) (X, \mathcal{B}, μ, T) is weakly mixing relative to $(X, \mathcal{B}_\gamma, \mu, T)$.

If $\mathcal{B}_\gamma = \mathcal{B}$, the system is said to be *distal*. In any case, \mathcal{B}_γ is called the *maximal distal factor*.

The exact sense in which Theorem 4.3 is nonconstructive will be addressed in the next section. What is striking about the structure theorem is that it has found direct application to finitary combinatorics, in a way I will now describe.

A k -coloring of the integers is simply a function c from \mathbb{Z} to $\{1, \dots, k\}$; think of $c(n)$ as the “color” assigned to the integer n . A set $A \subseteq \mathbb{Z}$ is *monochromatic* for the coloring if it is contained in $c^{-1}(i)$ for some i . Van der Waerden’s theorem [72] states the following:

⁹ With the obvious modifications, this argument works for nonseparable systems as well.

Theorem 4.4 *In any coloring of the integers with finitely many colors, there are arbitrarily long monochromatic arithmetic progressions.*

By a straightforward combinatorial “compactness” argument, this is equivalent to the following finitary version:

Theorem 4.5 *For every m and k there is an n large enough so that for any k coloring of the set $\{1, \dots, n\}$, there is a monochrome arithmetic progression of length m .*

If S is any set of integers, the *upper Banach density* of S is defined to be $\lim_n \sup_k |S \cap [k, n]|/n$. Thus the assertion that S has positive upper Banach density is equivalent to the assertion that for some $\delta > 0$, there are arbitrarily long intervals in the integers on which the set S has density at least δ . Szemerédi’s theorem is as follows:

Theorem 4.6 *Every set S of integers with positive upper Banach density has arbitrarily long arithmetic progressions.*

This is strictly stronger than van der Waerden’s theorem, since for any k coloring of the integers, there is a color i such that the set of elements S that are assigned color i has upper Banach density at least $1/k$. As was the case for van der Waerden’s theorem, Szemerédi’s theorem can also be stated in finitary terms:

Theorem 4.7 *For every k and $\delta > 0$, there is an n large enough, such that if S is any subset of $\{1, \dots, n\}$ with density at least δ , then S has an arithmetic progression of length k .*

Below, the least n satisfying the conclusion of the theorem for k and δ will be denoted $N_{SZ}(k, \delta)$.

Szemerédi provided a difficult combinatorial proof of Theorem 4.6 [68]. Soon after, however, Furstenberg [29] provided a new proof using ergodic-theoretic methods. His strategy was to recast the problem in measure-theoretic terms, by identifying sets of integers with elements of an appropriate measure preserving system. First, identify each set of integers S with its characteristic function χ_S , and view χ_S as an infinite binary sequence whose positions are indexed by the integers. We will define a T -invariant measure μ_S on the space $2^{\mathbb{Z}}$ of all such sequences, where Tx is the map which simply shifts each sequence to the left. If σ is a finite sequence of 0’s and 1’s, let $[\sigma]$ denote the set of elements x of $2^{\mathbb{Z}}$ that match σ starting at 0; it suffices to define $\mu_S([\sigma])$ for each σ . Pick a sequence I_j of intervals that witness the upper Banach density of S ; by thinning this sequence appropriately, one can arrange that for each σ , the density of occurrences of σ in each interval approaches a limit, r_σ . Define $\mu([\sigma]) = r_\sigma$. Thus, for each σ , $\mu([\sigma])$ is a measure of how often the pattern

σ occurs in the set S , and the statement that S has positive upper Banach density is equivalent to the statement that $\mu([1]) > 0$. When the following theorem is specialized to that measure space and the set A is taken to be the set $[1]$, Szemerédi’s theorem is an immediate consequence.

Theorem 4.8 *For any measure preserving system (X, \mathcal{B}, μ, T) any set A of positive measure, and any k there is an n such that*

$$\mu(A \cap T^{-n}A \cap T^{-2n}A \cap \dots \cap T^{(k-1)n}A) > 0.$$

Conversely, it is not hard to prove Theorem 4.8 from Szemerédi’s theorem. Thus the Furstenberg correspondence gives us a precise measure-theoretic analogue.

The structure theorem was a by-product of Furstenberg’s analysis, though his original proof of Szemerédi’s theorem [29] managed to avoid using the full strength of the structure theorem. Soon after, Furstenberg and Katznelson [31] presented a streamlined proof of an even stronger result, using the structure theorem in an essential way (see also [30,32]). It is easy to sketch the key ideas. If (X, \mathcal{B}, μ, T) is weak mixing, Theorem 4.8 holds for the following reason: since, on average, the events $T^{-in}A$ are close to uncorrelated for $i = 0, \dots, k - 1$, the measure of the set in question is close to $\mu(A)^k$ fairly often. In combinatorial terms, if the original set S is random enough, one would expect to find an arithmetic progression sooner or later, by dumb luck. If (X, \mathcal{B}, μ, T) is compact, Theorem 4.8 holds for an entirely different reason: for some n , the $T^{-n}A$ is guaranteed to return sufficiently close to A so that the intersection is nonempty after k iterations of T^{-n} . In combinatorial terms, sufficient regularity in the original set S is enough to guarantee the existence of an arithmetic progression.

For arbitrary spaces X , one formulates a slightly stronger inductive hypothesis; that is, one says what it means for a space (X, \mathcal{B}, μ, T) to be “SZ.” One then shows that every compact system is SZ, and that the property of being SZ is preserved under compact extensions, limits, and weakly mixing extensions. The Furstenberg structure theorem then implies that every space is SZ, which yields the desired conclusion.

5 Analysis of the structure theorem

It is commonly acknowledged that the ergodic theoretic proofs of Szemerédi’s theorem are nonconstructive. For example, Tao writes [69, page 2]:

This ergodic theory argument is the shortest and most flexible of all the

known proofs, and has been the most successful at leading to further generalizations of Szemerédi’s theorem. . . On the other hand, the infinitary nature of the argument means that it does not obviously provide any effective bounds for the quantity $N_{SZ}(k, \delta)$.

But mathematicians are somewhat vague as to the precise source of the non-constructivity. In surveying the background to the recent Tao-Green proof that there are arbitrarily long arithmetic progressions in the primes [40], Kra writes:

Furstenberg’s proof relies on a compactness argument, making it difficult to extract any explicit bounds in the finite version of Szemerédi’s theorem. [53, page 7]

Kra is referring to the compactness argument that is implicit in the Furstenberg correspondence principle, namely, the iterative thinning of an infinite sequence of intervals witnessing the upper Banach density of S used in the construction of the measure μ . Furstenberg himself writes, in his *Mathematical Review* of Gowers’s elementary proof of Szemerédi’s theorem [39]:

However, the ergodic-theoretic approach depends essentially on passing to a limit whereby a set $\{1, 2, 3, \dots, N\}$ is replaced by a measure space, and the translations $n \rightarrow n + a$ are replaced by measure preserving transformations of this space. In passing to this limit one loses sight of the size N of the interval $\{1, 2, 3, \dots, N\}$. As a result this approach is incapable of giving any information regarding $[N_{SZ}(k, \delta)]$ beyond the fact that it is finite.

But this use of nonconstructivity is fairly mild. The ergodic-theoretic arguments show that for any measurable space (X, \mathcal{B}) , any measure μ on that space, any set A in \mathcal{B} , any k , and any $\delta > 0$, there is an n such that if $\mu(A) \geq \delta$, $\mu(\bigcap_{i < k} T^{-in} A) > 0$. In particular, this holds for the fixed space (X, \mathcal{B}) and the fixed set A used in the Furstenberg correspondence. Now notice that the set of probability measures μ on (X, \mathcal{B}) is compact in the weak-* topology, as is the closed subset of measures μ satisfying $\mu(A) \geq \delta$. Moreover, fixing k and δ , the function mapping μ to the least n such that $\mu(\bigcap_{i < k} T^{-in} A) > 0$ is continuous in the weak-* topology (with the discrete topology on \mathbb{N}), since the conclusion involves only finitely many values of μ . Thus, for every k and δ , there is a bound $n(k, \delta)$ that is independent of μ . It is a straightforward combinatorial exercise to translate this to a bound on N_{SZ} .¹⁰

¹⁰ Specifically, given k and δ , let $n = n(2k, \delta)$, and let $l = 2(k - 1)n + 1$. Then if S is any subset of $[1, l]$ of density at least δ , S must have an arithmetic progression of length k , and, moreover, one with common difference at most n . To see this, note that otherwise, laying copies of S side by side, we obtain arbitrarily large sets of density at least δ with no arithmetic progression of length $2k$ and common difference at most n ; but then any measure obtained from such a sequence via the

In formal logical terms, this application of compactness can be reduced to an appeal to a principle known as “weak König’s lemma,” which asserts that any infinite, finitely branching tree has an infinite path. This principle is, indeed, nonconstructive; an argument due to Kleene shows that such a path cannot always be computed, even if the tree is effectively presented. But the Jockusch-Soare “low basis theorem” [46] guarantees that there is always a path of low complexity. Building on a seminal conservation result due to Harvey Friedman, proof theoretic research has provided a number of ways of eliminating the use of weak König’s lemma from proofs of combinatorial statements.¹¹ Moreover, modern proof mining methods [51] make it possible to do this effectively in practice. In short, if the Furstenberg correspondence principle were the only nonconstructive feature of the argument, it would not be hard to reinterpret the proof in computational or combinatorial terms.

Of course, the transfinite iteration involved in the structure theorem should seem suspect. But, from a constructive point of view, there is nothing inherently wrong with a definition by transfinite recursion. Indeed, many axiomatizations of constructive mathematics allow induction and recursion on inductively defined sets. To illustrate, define a tree T on \mathbb{N} to be a set of finite sequences of natural numbers closed under initial segments. Think of each sequence as providing an “address” of the node; thus $()$ is the root, and any immediate children of an element σ are the elements of the form $\sigma^\frown(n)$. Such a tree is *well-founded* if there is no infinite path; that is, for every function $f : \mathbb{N} \rightarrow \mathbb{N}$ there is a natural number n such that the sequence $(f(0), f(1), f(2), \dots, f(n))$ has left the tree. It will be convenient to restrict our attention to trees that are *full*, so that any node σ in the tree either has no children, or $\sigma^\frown(n)$ is in the tree for every n . Let e denote the tree $\{()\}$ with just one node, and if σ is any node of T , let T_σ denote the subtree $\{\tau \mid \sigma^\frown\tau \in T\}$ rooted at σ . Classically, one can show that the set W of full well-founded trees on \mathbb{N} can be generated by the following two clauses:

- e is in W ; and
- If $f : \mathbb{N} \rightarrow W$ is any sequence elements of W , and T is the tree such that for every n the n th subtree $T_{(n)}$ is equal to $f(n)$, then T is in W .

From a constructive point of view, the two characterizations of the set of well-founded trees are not equivalent, and the latter, inductive, definition is preferred. With that characterization, one can justify the following principle of recursion: one can specify a function F from the set of well-founded trees

Furstenberg correspondence fails to have the property guaranteed by our choice of n . (Alternatively, one can let $n = n(k, \delta')$ for a slightly smaller value of δ' , pick l large enough, and then remove a small number of elements from the end of S to eliminate “wraparound” effects.)

¹¹ There is not enough space here for me to survey what is known about weak König’s lemma; but see [4, footnote 12] for an overview and references.

to any other set X with two clauses,

$$\begin{aligned} F(e) &= a \\ F(T) &= G(\lambda n F(T_n)) \quad \text{if } T \text{ is not } e, \end{aligned} \tag{3}$$

where a is an element of X and G is a function from sequences of trees to X .

Georg Kreisel's theory ID_I [57,17] provides an axiomatic basis for reasoning about such inductive definitions. Take classical first-order Peano arithmetic, PA , to be formulated in a language with symbols for each primitive recursive function and relation. The axioms of PA include basic axioms defining these functions and relations, and the schema of induction:

$$\varphi(0) \wedge \forall x (\varphi(x) \rightarrow \varphi(x+1)) \rightarrow \forall x \varphi(x),$$

where φ is any formula in the language. Heyting arithmetic, HA , is the analogous theory, founded on intuitionistic first-order logic. ID_I is an extension of PA with additional predicates P intended to denote the fixed-points of certain types of inductive definitions. Specifically, let $\psi(P, x)$ denote a formula in the language of arithmetic with one additional predicate or set symbol P that has only positive occurrences in ψ . Formally, positivity means that the expressions $P(t)$ occur unnegated when the definition is written in negation-normal form; intuitively, this means that the definition can only use positive information as to which elements satisfy P . This determines a monotone operator Γ_ψ from sets to sets defined by

$$\Gamma_\psi(S) = \{x \in \mathbb{N} \mid \psi(S, x)\}.$$

Such an operator has a least fixed point, $I = \bigcap \{S \mid \Gamma_\psi(S) \subseteq S\}$. The theory ID_I adds a new predicate symbol P for each such ψ , intended to denote this fixed point, together with the following axioms:

- $\forall x (\psi(P, x) \rightarrow P(x))$
- $\forall x (\psi(\theta/P, x) \rightarrow \theta(x)) \rightarrow \forall x (P(x) \rightarrow \theta(x))$, for each formula θ .

Here the notation $\psi(\theta/P)$ denotes the result of replacing each atomic formula $P(t)$ with $\theta(t)$, renaming bound variables to prevent collisions. The first axiom implies that P is closed with respect to Γ_ψ , while the second axiom schema expresses that P is the smallest set closed under Γ_ψ , at least, to the extent that it is possible to do so within a first-order language.

One can also design theories of inductive definitions based on intuitionistic logic. In this case, however, one needs to be more careful in specifying the positivity requirement on ψ . One option is to insist that P does not occur in the antecedent of any implication, where $\neg\eta$ is taken to abbreviate $\eta \rightarrow \perp$. Such a definition is said to be *strictly positive*, and we can denote the corresponding axiomatic theory $ID_I^{i,sp}$. An even more restrictive requirement is to insist that

θ is of the form $\forall y (y \prec x \rightarrow P(y))$, where \prec is a primitive recursive relation. These are called *accessibility* inductive definitions, and serve to pick out the well-founded part of the relation. In the case where \prec is the “child-of” relation on a tree, the inductive definition picks out the well-founded part of that tree. We will denote the corresponding theory $ID_1^{i,acc}$. The point is this: in $ID_1^{i,sp}$, one can define the set of computable well-founded trees and justify the principle of recursion on those trees. If one restricts attention to primitive recursive trees, the same goes through in $ID_1^{i,acc}$. Moreover, these theories have natural computational interpretations, and are commonly accepted as constructively valid.

What, then, makes the Furstenberg proof nonconstructive? The answer is found where mathematicians are unlikely to expect it, namely, in the fairly mundane use of limits, or projections, in the argument. Lemma 4.1 and its relativized version 4.2 make use of the ergodic theorem, and we have already seen, in Section 3, that the theorem does not admit a direct computable interpretation.¹² The transfinite iteration then amplifies the problem, yielding a transfinite sequence of nonconstructive definitions.

At this point, the methods of descriptive set theory and effective descriptive set theory are helpful in characterizing the complexity of the resulting structures. Beznay and Foreman [10] have shown that the Furstenberg construction can exhaust the countable ordinals, in the following sense. Define the *order* of a measure preserving system to be the length of the shortest tower satisfying the conclusion of the Furstenberg structure theorem. Beznay and Foreman have shown:

Theorem 5.1 *The set of orders of measure preserving systems is exactly the set of countable ordinals.*¹³

In fact, if Y is a factor of a measure preserving system \mathcal{X} , there is a *largest* relatively compact extension $Z(Y)$ of Y , and the shortest Furstenberg tower is obtained by taking this extension at every stage. Since we are assuming \mathcal{X} is separable, any element of $L^2(\mathcal{X})$ can be coded as a set of natural numbers. It is not hard to show, as Beznay and Foreman do, that $Z(Y)$, viewed as a set of elements of $L^2(\mathcal{X})$, can then be defined by an arithmetic formula in \mathcal{X} and Y . Towsner and I have shown that there is a coding of the factors themselves as sets of natural numbers such that $Z(Y)$ has a Δ_2 arithmetic definition in \mathcal{X} and Y , and, moreover, the map $Y \mapsto Z(Y)$ is monotone. With this coding,

¹²The argument’s use of projections onto factors also requires arithmetic comprehension; see [7].

¹³Beznay and Foreman provide, moreover, a Borel construction that assigns to any countable linear ordering a separable measure preserving system whose order is the well-founded part, showing that the collection of measure distal systems is a complete Π_1^1 set under Borel reducibility.

the maximal distal factor is therefore an instance of a monotone arithmetic inductive definition. Seminal results due to Spector [67] together with “stage comparison” methods due to Moschovakis, Aczel, and Kunen show that such inductive definitions terminate by the Church-Kleene ordinal $\omega_1^{CK,X}$, that is, the least ordinal that is not computable relative to a code X for the original system (see [60] for details). Thus we have:

Theorem 5.2 *Let X code any measure preserving system. Then the height of the Furstenberg tower is less than or equal to $\omega_1^{CK,X}$. For any $\alpha < \omega_1^{CK,X}$, the α th factor can be computed from $H_{2,\alpha}^X$, i.e. the $2 \cdot \alpha$ th hyperarithmetic set relative to X .¹⁴*

Note that if α is a limit, $2 \cdot \alpha = \alpha$. Towsner and I suspect that the complexity lower bounds given by Theorem 5.2 are sharp, at least for limit ordinals, in the sense that for every set A there is a measure preserving system computable from A such that for every computable ordinal α , H_α is computable from the α th factor. Proving such a theorem will require a careful and subtle analysis, but a cruder analysis, based on the methods of Belezny and Foreman, shows that in the sense of reverse mathematics (see [66]), the structure theorem is axiomatically strong:

Theorem 5.3 *Over ACA_0 , the Furstenberg structure theorem is equivalent to the Π_1^1 comprehension axiom.*

Thus the Furstenberg tower associated to an arbitrary measure preserving system may be a wildly uncomputable object. And yet, references to this object allow us to prove a finitary combinatorial statement with explicit computational content. This state of affairs calls for metamathematical explanation: we wish to understand how this detour through the infinite works, and the extent to which it can be reconciled with a computational view of mathematics.

Towsner and I offer a two-part explanation. The first part is an analysis of the Furstenberg proof in axiomatic terms. We have shown that the definition of $Z(Y)$ in terms of Y can be given by a *positive* Σ_4 arithmetic formula, which implies:

Theorem 5.4 *Let X code any measure preserving system. Then the code \hat{Y} for the maximal distal factor of X has a positive arithmetic inductive definition relative to X .*

This enables us to develop a version of the structure theorem in ID_1 . With careful attention to the other analytic and combinatorial objects involved, we are then able to show:

¹⁴See Ash and Knight [2] for an introduction to hyperarithmetic set theory.

Theorem 5.5 *The Furstenberg proof of Szemerédi’s theorem can be carried out in ID_1 .*

The second part of our analysis shows that proofs of statements like Szemerédi’s theorem in ID_1 can always be interpreted in constructive terms. The next section is devoted to filling out this claim.

6 The constructive content of ID_1

In the logical terminology, a Π_2 sentence is one of the form $\forall \bar{x} \exists \bar{y} R(\bar{x}, \bar{y})$, where \bar{x} and \bar{y} are tuples of variables ranging over the natural numbers, and R is a primitive recursive relation.¹⁵ Any such sentence can be understood as making the computational assertion that an algorithm which, on input \bar{x} , searches for a tuple \bar{y} satisfying R is bound to terminate. In particular, the finitary statements of van der Waerden’s theorem and Szemerédi’s theorem in Section 4 have that form.

One way of assessing the constructive content of a theory, T , is to characterize its Π_2 consequences. Here are two forms such a characterization can have:

- (1) Every Π_2 sentence provable in T is also provable in a constructive theory, T' .
- (2) Every Π_2 sentence provable in T is witnessed by an elements of a particular class of computable functions, C .

The two types of results are often closely related: the relevant constructive theory, T' , is often based on principles that reflect natural programming constructs, and, conversely, a natural characterization of the relevant functions can often be read off straightforwardly from T . These patterns leave the choice of T' and C open, and the extent to which a particular choice of T' or a particular characterization of C is illuminating is subject to debate. The goal of this section is present an informative characterization the constructive content of ID_1 in these terms. To set the stage, however, it will be helpful to review analogous facts regarding the constructive content of PA . The first result states that PA is a conservative extension of HA for Π_2 sentences:

Theorem 6.1 *Every Π_2 sentence provable in PA is provable in HA .*

¹⁵It is well known that the expressive power of a Π_2 formula does not change if one replaces “primitive recursive” with “computable” or “ Δ_0 definable,” etc. From an axiomatic standpoint, “primitive recursive” is a reasonable stand-in for “straightforwardly computable.”

To characterize the class of functions that suffice to witness the Π_2 theorems of PA , define the set of *finite types* inductively, as follows:

- N is a finite type; and
- assuming σ and τ are finite types, so are $\sigma \times \tau$ and $\sigma \rightarrow \tau$.

In the “full” set-theoretic interpretation, N denotes the set of natural numbers, $\sigma \times \tau$ denotes the set of ordered pairs consisting of an element of σ and an element of τ , and $\sigma \rightarrow \tau$ denotes the set of functions from σ to τ . But we can also view the finite types as nothing more than datatype specifications of computational objects. The set of *primitive recursive functionals of finite type* is a set of computable functionals obtained from the use of explicit definition (λ abstraction), application, pairing, and projections, and a scheme of primitive recursion:

$$\begin{aligned} F(0) &= a \\ F(n+1) &= G(n, F(n)) \end{aligned}$$

Here, the range of F may be any finite type.

Theorem 6.2 *Every Π_2 theorem of PA is witnessed by a primitive recursive functional of type $N^k \rightarrow N$.*

There are two principal ways of obtaining the pair of results we have just described. In both cases, the first step is to use the Gödel-Gentzen double-negation translation to interpret PA in HA . The interpretation does not, unfortunately, preserve Π_2 sentences, since $\forall \bar{x} \exists \bar{y} R(\bar{x}, \bar{y})$ is interpreted as $\forall \bar{x} \neg \neg \exists \bar{y} R(\bar{x}, \bar{y})$. From there, one has two choices.

- (1) Use the double-negation interpretation to interpret PA in HA . Use the Friedman-Dragalin A-translation [22,26] to “repair” the interpretation of Π_2 sentences to yield the desired conservation result.¹⁶ Then use Kreisel’s modified version of Kleene’s realizability to extract a witness (see [54,71]).
- (2) Use the Dialectica interpretation [37,5] to extract a primitive recursive functional witnessing the conclusion. This last step can also be interpreted in HA , yielding the conservation result.

Both methods establish the two theorems, and are equally good, from that perspective. But now one can ask, what happens when one applies the results to particular proofs? The first method has been used to extract interesting algorithms from proofs of classical results, such as Dickson’s lemma [11]. But in ongoing research in “proof mining,” the experience of Kohlenbach and his students has shown that the Dialectica interpretation is generally a more powerful and effective tool.¹⁷

¹⁶ An alternative method of repairing Π_2 sentences is described in [3,20].

¹⁷ This experience is born out by an MS thesis [41] written by a student of mine,

Let us now lift these results to ID_1 . We have already met the appropriate intuitionistic counterpart, in Section 6:

Theorem 6.3 *Every Π_2 sentence provable in ID_1 is provable in $ID_1^{i,acc}$.*

One can characterize functions witnessing the Π_2 theorems of ID_1 with a natural extension of the primitive recursive functionals of finite type, described in [5, Section 9.1]. We simply extend the finite types by adding a new base type, Ω , which is intended to denote the set of well-founded (full) trees on the natural numbers. We add a constant, e , which denote the tree with just one node, and two new operations: Sup , of type $(N \rightarrow \Omega) \rightarrow \Omega$, which forms a new tree from a sequence of subtrees, and Sup^{-1} , of type $\Omega \rightarrow (N \rightarrow \Omega)$, which returns the immediate subtrees of a nontrivial tree. Finally, we simply add the principle of recursive definition corresponding to (3):

$$\begin{aligned} F(e) &= a \\ F(Sup(h)) &= G(\lambda n F(h(n))), \end{aligned}$$

where the range of F can be any of the new types. Call these the *primitive recursive tree functionals*.

Theorem 6.4 *Every Π_2 theorem of ID_1 is witnessed by a primitive recursive tree functional of type $N^k \rightarrow N$.*

As was the case with Peano arithmetic, there are two distinct ways of arriving at Theorems 6.3 and 6.4. The first involves using a combination of the double-negation translation and a complex forcing relation due to Buchholz [16,1] to prove Theorem 6.3, after which modified realizability provides Theorem 6.4. Until recently, this was the only way of obtaining these results, short of passing through an ordinal analysis of ID_1 . In particular, there was no way of obtaining these two theorems using a variant of the Dialectica interpretation, a disappointing fact that is highlighted in [5, Section 9.8]. Towsner and I [8] have now closed the gap by providing a Dialectica interpretation of ID_1 that is clean and remarkably simple. We expect that this translation will prove to be a valuable tool in the analysis of proofs, like the ergodic theoretic proof of Szemerédi's theorem, that rely on inductively defined sets and structures.

Aaron Hertz, who applied the Dialectica interpretation to obtain a constructive proof of the Hilbert basis theorem (subsuming Dickson's lemma). Whereas Berger et al. [11] had difficulty with more than two variables, Hertz easily obtained terms witnessing the Dialectica interpretation of the full version.

7 Conclusion

The results I have described provide a strategy for obtaining a purely combinatorial version of the Furstenberg proof of Szemerédi’s theorem: formalize the proof in ID_1 , and then apply our Dialectica translation. Of course, the final goal is not to obtain a formal derivation, but, rather, an explicit combinatorial proof that can be read and appreciated in ordinary mathematical terms. Henry Towsner and I are currently working on obtaining such a proof. Our analysis is thus similar to Girard’s “unwinding” [35] of a topological dynamical proof of van der Waerden’s theorem by Furstenberg and Weiss [33], only more involved.

In 1998, Gowers was awarded a Fields Medal, in part, for his use of Fourier analytic methods (and combinatorial results due to Freiman) to obtain elementary bounds on Szemerédi’s theorem [39]; that is, bounds in terms of a fixed iterate of the exponential function. Gowers and Tao have both surmised that a careful analysis of Szemerédi’s original proof will show that it yields primitive recursive bounds (assuming one makes use of elementary bounds on the Hales-Jewett theorem obtained by Shelah). Tao has [69] presented another combinatorial proof of Szemerédi’s theorem inspired by Furstenberg’s proof, and remarked, in passing:

It may be possible in principle to extract some bound for $N_{SZ}(k, \delta)$ directly from [Furstenberg’s] original argument via proof theory, using such tools as Herbrand’s theorem.

The bounds resulting from Tao’s proof [69] seem to be slightly worse than Ackermannian.

A by-product of the program that Towsner and I are pursuing would, indeed, be a bound on the rate of growth of N_{SZ} expressed in terms of the functional calculus described in the last section. But functions in that calculus can have astronomical rates of growth, and we do not expect that the expression we extract from Furstenberg’s proof will yield useful bounds, without additional work. There are, nonetheless, good reasons for pursuing the program we have set. Our goal is to obtain a perspicuous new proof of Szemerédi’s theorem, one that will clarify the combinatorial essence of the Furstenberg approach and yield new combinatorial ideas and methods. The work may, for example, lead to interesting generalizations and variants of Szemerédi’s theorem. It may also point the way to finding combinatorial theorems that require the full logical strength of ID_1 , akin to similar combinatorial independences obtained by Harvey Friedman (see, for example, [27,28]).

From a metamathematical perspective, our analysis is also interesting in its own right. The fact that abstract, infinitary methods can have direct bearing

on finitary concerns is a striking phenomenon, and one that should be explored. Understanding how this works in the case of ergodic theory and combinatorics is an important component of the more general project of understanding of the role that infinitary methods play in mathematics, and the ways that modern methods can be understood in computational terms.

References

- [1] Klaus Aehlig. *On fragments of analysis with strengths of finitely iterated inductive definitions*. PhD thesis, University of Munich, 2003.
- [2] C. J. Ash and J. Knight. *Computable structures and the hyperarithmetical hierarchy*. North-Holland Publishing Co., Amsterdam, 2000.
- [3] Jeremy Avigad. Interpreting classical theories in constructive ones. *J. Symbolic Logic*, 65:1785–1812, 2000.
- [4] Jeremy Avigad. Number theory and elementary arithmetic. *Philosophia Mathematica*, 11:257–284, 2003.
- [5] Jeremy Avigad and Solomon Feferman. Gödel’s functional (“Dialectica”) interpretation. In [19], pages 337–405.
- [6] Jeremy Avigad, Philipp Gerhardy, and Henry Towsner. Local stability of ergodic averages. To appear in *Trans. Am. Math. Soc.*
- [7] Jeremy Avigad and Ksenija Simic. Fundamental notions of analysis in subsystems of second-order arithmetic. *Ann. Pure Appl. Logic*, 139:138–184, 2006.
- [8] Jeremy Avigad and Henry Towsner. Functional interpretation and inductive definitions. arXiv:0802.1938.
- [9] June Barrow-Green. *Poincaré and the three body problem*. American Mathematical Society, Providence, RI, 1997.
- [10] Ferenc Beleznyai and Matthew Foreman. The complexity of the collection of measure-distal transformations. *Ergod. Thy. Dynam. Sys.*, 16:929–962, 1996.
- [11] Ulrich Berger, Helmut Schwichtenberg, and Monika Seisenberger. The Warshall algorithm and Dickson’s lemma: two examples of realistic program extraction. *J. Automated Reasoning*, 26:205–221, 2001.
- [12] George D. Birkhoff. Proof of the ergodic theorem. *Proc. Nat. Acad. Sci.*, 17:656–660, 1931.
- [13] Errett Bishop. An upcrossing inequality with applications. *Michigan Math. J.*, 13:1–13, 1966.

- [14] Errett Bishop. *Foundations of Constructive Analysis*. McGraw-Hill, New York, 1967.
- [15] Errett Bishop. A constructive ergodic theorem. *J. Math. Mech.*, 17:631–639, 1967/1968.
- [16] Wilfried Buchholz. The $\Omega_{\mu+1}$ -rule. In [17], pages 188–233. 1981.
- [17] Wilfried Buchholz, Solomon Feferman, Wolfram Pohlers, and Wilfried Sieg. *Iterated inductive definitions and subsystems of analysis: recent proof-theoretical studies*. Springer, Berlin, 1981.
- [18] John Burgess. From the outside in. Preprint, based on a lecture delivered to the ASL Spring Meeting, Montreal, 2006.
- [19] Samuel R. Buss, editor. *The Handbook of Proof Theory*. North-Holland, Amsterdam, 1998.
- [20] Thierry Coquand and Martin Hofmann. A new method of establishing conservativity of classical systems over their intuitionistic version. *Mathematical Structures in Computer Science*, 9:323–333, 1999.
- [21] Richard Dedekind. *Sur la théorie des nombres entiers algébriques*. Gauthier-Villars, Paris, 1877. Translated as *Theory of algebraic integers* with an editorial introduction by John Stillwell, Cambridge University Press, Cambridge, 1996.
- [22] Albert Dragalin. *Mathematical intuitionism: introduction to proof theory*. American Mathematical Society, 1988.
- [23] William Ewald, editor. *From Kant to Hilbert: a source book in the foundations of mathematics*. Clarendon Press, Oxford, 1996. Volumes 1 and 2.
- [24] Solomon Feferman. Hilbert’s program relativized: proof-theoretical and foundational reductions. *J. Symbolic Logic*, 53:364–384, 1988.
- [25] Solomon Feferman. Relationships between constructive, predicative, and classical systems of analysis. In Vincent F. Hendricks, Stig Andur Pederson, and Klaus Frovin Jorgensen, editors, *Proof Theory: History and Philosophical Significance*, pages 221–236. Kluwer Academic Publishers, 2000.
- [26] Harvey M. Friedman. Classically and intuitionistically provable functions. In H. Müller and D. Scott, editors, *Higher Set Theory*, pages 21–27. Springer, Berlin, 1978.
- [27] Harvey M. Friedman. Finite functions and the necessary use of large cardinals. *Annals of Mathematics*, 148:803–893, 1998.
- [28] Harvey M. Friedman. Internal finite tree embeddings. In *Reflections on the foundations of mathematics, 1998*, pages 60–91. Assoc. Symbol. Logic, Urbana, IL, 2002.
- [29] H. Furstenberg. Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions. *J. Analyse Math.*, 31:204–256, 1977.

- [30] H. Furstenberg. *Recurrence in ergodic theory and combinatorial number theory*. Princeton University Press, Princeton, N.J., 1981.
- [31] H. Furstenberg and Y. Katznelson. An ergodic Szemerédi theorem for commuting transformations. *J. Analyse Math.*, 34:275–291, 1979.
- [32] H. Furstenberg, Y. Katznelson, and D. Ornstein. The ergodic theoretical proof of Szemerédi’s theorem. *Bull. Amer. Math. Soc. (N.S.)*, 7:527–552, 1982.
- [33] H. Furstenberg and B. Weiss. Topological dynamics and combinatorial number theory. *J. Analyse Math.*, 34:61–85, 1978.
- [34] Philipp Gerhardy and Ulrich Kohlenbach. General logical metatheorems for functional analysis, 2006. To appear in *Trans. Am. Math. Soc.*
- [35] Jean-Yves Girard. *Proof theory and logical complexity*. Bibliopolis, Naples, 1987.
- [36] Kurt Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatsh. Math. Phys.*, 38(1):173–198, 1931. Reprinted with English translation in [38], volume 1, pages 144–195.
- [37] Kurt Gödel. Über eine bisher noch nicht benützte Erweiterung des finiten Standpunktes. *Dialectica*, 12:280–287, 1958. Reprinted with English translation in [38], volume II, pages 241–251.
- [38] Kurt Gödel. *Collected works*. Oxford University Press, New York, 1986–2003. Edited by Solomon Feferman et al. Volumes I–V.
- [39] W. T. Gowers. A new proof of Szemerédi’s theorem. *Geom. Funct. Anal.*, 11:465–588, 2001.
- [40] Ben Green and Terence Tao. The primes contain arbitrarily long arithmetic progressions. *Ann. of Math.*, to appear.
- [41] Aaron Hertz. A constructive version of the Hilbert basis theorem. Master’s thesis, Carnegie Mellon University, 2004.
- [42] David Hilbert. Neubegründung der Mathematik. *Abhandlungen aus dem mathematischen Seminar der Hamburgischen Universität*, 1:157–177, 1922. Translated by William Ewald in [23], volume 2, pages 1115–1134.
- [43] David Hilbert. Die logischen Grundlagen der Mathematik. *Mathematische Annalen*, 88:151–165, 1923. Translated by William Ewald in [23], volume 2, pages 1134–1148.
- [44] Michael Hochman. Upcrossing inequalities for stationary sequences and applications. arXiv:math/0608311v2.
- [45] V. V. Ivanov. Oscillations of averages in the ergodic theorem. *Dokl. Akad. Nauk*, 347:736–738, 1996.
- [46] Carl Jockusch and Robert Soare. Π_1^0 classes and degrees of theories. *Trans. Am. Math. Soc.*, 173:33–56, 1972.

- [47] Roger L. Jones, Robert Kaufman, Joseph M. Rosenblatt, and Máté Wierdl. Oscillation in ergodic theory. *Ergod. Thy. Dynam. Sys.*, 18:889–935, 1998.
- [48] Roger L. Jones, Joseph M. Rosenblatt, and Máté Wierdl. Counting in ergodic theory. *Canad. J. Math.*, 51:996–1019, 1999.
- [49] A. G. Kachurovskii. Rates of convergence in ergodic theorems. *Uspekhi Mat. Nauk*, 51:73–124, 1996. Translation in Russian Math. Surveys 51: 653–703, 1996.
- [50] Steven Kalikow and Benjamin Weiss. Fluctuations of ergodic averages. *Illinois J. Math.*, 43:480–488, 1999.
- [51] Ulrich Kohlenbach. *Applied proof theory: proof interpretations and their use in mathematics*. Springer, Berlin, to appear.
- [52] B. O. Koopman and J. von Neumann. Dynamical systems of continuous spectra. *Proc. Nat. Acad. Sci. USA*, 18:255–263, 1932.
- [53] Bryna Kra. The Green-Tao theorem on arithmetic progressions in the primes: an ergodic point of view. *Bull. Amer. Math. Soc. (N.S.)*, 43:3–23 (electronic), 2006.
- [54] Georg Kreisel. On weak completeness of intuitionistic predicate logic. *J. Symbolic Logic*, 27:139–158, 1962.
- [55] Georg Kreisel. On the interpretation of non-finitist proofs, part I. *J. Symbolic Logic*, 16:241–267, 1951.
- [56] Georg Kreisel. Interpretation of analysis by means of constructive functionals of finite type. In Arendt Heyting, editor, *Constructivity in Mathematics*, pages 101–128. North-Holland, Amsterdam, 1959.
- [57] Georg Kreisel. Generalized inductive definitions. Stanford Report on the Foundations of Analysis (mimeographed), 1963.
- [58] Ulrich Krengel. On the speed of convergence in the ergodic theorem. *Monatsh. Math.*, 86:3–6, 1978/79.
- [59] Pierre-Simon Laplace. *Essai philosophique sur les probabilités*. Courcier, Paris, 1814. Fifth (1825) edition translated by Andrew I. Dale as *Philosophical essay on probabilities*, Springer, New York 1995.
- [60] Yiannis Moschovakis. *Elementary induction on abstract structures*. North-Holland, Amsterdam, 1974.
- [61] Henri Poincaré. Mémoire sur les courbes définies par une équation différentielle. *Journal de Mathématiques*, 7:375–422, 1881.
- [62] Henri Poincaré. *Science et Méthode*. Flammarion, Paris, 1908. Translated by Francis Maitland as *Science and method*, Dover Publications, New York, 1952.
- [63] Marian B. Pour-El and J. Ian Richards. *Computability in analysis and physics*. Springer-Verlag, Berlin, 1989.

- [64] Frederick Riesz. Another proof of the mean ergodic theorem. *Acta Univ. Szeged. Sect. Sci. Math.*, 10:75–76, 1941.
- [65] Ksenija Simic. The pointwise ergodic theorem in subsystems of second-order arithmetic. *J. Symbolic Logic*, 72:45–66, 2007.
- [66] Stephen G. Simpson. *Subsystems of Second-Order Arithmetic*. Springer, Berlin, 1999.
- [67] Clifford Spector. Inductively defined sets of natural numbers. In *Infinitistic methods*, Pergamon, New York, 1961, 97–102.
- [68] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith.*, 27:199–245, 1975.
- [69] Terence Tao. A quantitative ergodic theory proof of Szemerédi’s theorem. *Electron. J. Combin.*, 13:Research Paper 99, 49 pp. (electronic), 2006.
- [70] Terence Tao. Norm convergence of multiple ergodic averages for commuting transformations. To appear in *Ergod. Thy. Dynam. Sys.* arXiv:0707.1117.
- [71] A. S. Troelstra. Realizability. In [19], pages 407–473.
- [72] B. L. van der Waerden. Beweis einer Baudetschen Vermutung. *Nieuw. Arch. Wisk.*, 15: 212–216, 1927.
- [73] John von Neumann. Proof of the quasi-ergodic hypothesis. *Proc. Nat. Acad. Sci.*, 18:70–82, 1932.
- [74] V. V. V’yugin. Ergodic theorems for individual random sequences. *Theor. Comp. Sci.*, 207:343–361, 1998.
- [75] Klaus Weihrauch. *Computable analysis: an introduction*. Springer-Verlag, Berlin, 2000.
- [76] Richard S. Westfall. *Never at rest: a biography of Isaac Newton*. Cambridge University Press, Cambridge, 1980.