

# Realism in Mathematics: The Case of the Hyperreals

Kenny Easwaran and Henry Towsner

January 30, 2018

## 1 Introduction

A traditional question in the philosophy of mathematics is whether abstract mathematical objects really exist, just as planets and atoms and giraffes do, independently of the human mind. Views on this question are often grouped into “platonist”, “nominalist”, and “constructivist” groups depending on whether they state that mathematical objects have their own independent existence, or don’t truly exist at all, or exist but only as constructions of the human mind. This paper will not directly address this question. We expect that to the extent that each of these views is defensible, they will not settle the questions that we are interested in (though there may be some affinities between some possible answers to these questions). Rather, we will focus on two other questions of mathematical realism: the question of which mathematical claims can be taken as meaningful and true within mathematics, and which mathematical ideas can be taken as applying to the physical world as part of a scientific theory.

We will focus on these questions for the case of the “hyperreals” of non-standard analysis, introduced by Abraham Robinson in the 1960’s (which we will describe in greater detail later). Because of their mathematical properties, the hyperreals have attracted attacks and defenses in ways that are uncommon for mathematical entities. Our central claim is that many of the views on both sides of these debates conflate two distinct philosophical questions. Defenders of the hyperreals give realist answers to both questions, while attackers give antirealist answers to both. We defend a realist answer to one question and an antirealist answer to the other. We think the distinction between these two questions is important not just for thinking about the hyperreals, but for other related mathematical entities, and particularly those whose existence depend on the Axiom of Choice.

### 1.1 Factualism

The first question about mathematical realism that we are interested in is the factual question of whether a particular claim is true or false within mathematics. This question is meant to be understood in a sense that is neutral between

the different ontological views of mathematics. The nominalist, platonist, and constructivist can all agree that the statement “ $2+2 = 4$ ” has one status, which we can call “truth” (even though some of these philosophical views might deny that it is truth in the same sense that applies to statements about the physical world), while the statement “ $2+2 = 5$ ” has a different status, which we can call “falsehood”. Many platonists (and some defenders of other ontological views of mathematics) claim that *every* well-formed mathematical statement has a definite truth value of this sort — this view is sometimes called “truth-value realism”. [Linnebo, 2013, Section 1.4]

Platonists usually claim that Cantor’s continuum hypothesis has one of these two statuses (though they don’t necessarily agree which status it has, or whether we have any information relevant to this question), but many nominalists and constructivists deny (or at least refuse to assert) that the continuum hypothesis has one or the other status.

We claim that regarding this sort of question, whether or not all well-formed mathematical statements have a definite fact, there is a fact about the Axiom of Choice, and the fact is that it is true, as well as all of its consequences. Among these consequences are the existence of hyperreals, as well as various mathematical and physical facts that the hyperreals are sometimes used to calculate.

## 1.2 Applicability

The other question of realism that we focus on is more closely connected to realism in scientific theorizing than to questions squarely in philosophy of mathematics itself. [Chakravartty, 2017] The question is about the extent to which the theories we build using various purported mathematical entities accurately describe the world, or are merely instrumentally useful in modeling the world and making observational predictions. As a simple example, the structure of the natural numbers seems to accurately represent counts of predators and prey (at least, when talking about mammals or birds or other macroscopic animals with clear individuation). But while the real-valued differential equations of the Lotka-Volterra model can often be useful in predicting or understanding the ways the predator and prey populations will change, the infinite precision of various non-integer values that show up are not taken to represent the actual numbers of predators or prey in the ecosystem. It is our contention that all applications of the hyperreals to the sciences — whether in physics, biology, cognitive science, or other — are merely modeling instruments like this particular use of the real numbers.

The hyperreals are perfectly good mathematical entities, but we still need to be careful when using them to theorize about the physical world.

## 2 The hyperreals

The hyperreals are an extension of the real numbers to include infinitesimal numbers and, necessarily, also infinite real numbers. Famously, these extensions

were the basis of the theory of non-standard analysis developed by Abraham Robinson in the 1960's. (See [Robinson, 1996], first published in 1966.)

Specifically, the hyperreals are *elementarily equivalent* to the reals—they satisfy the exact same sentences of first-order logic. To make this precise, we must decide what the language is—that is, which operations and relations are present. Since we are talking about extensions of the reals, we assume we are at least considering the reals as an ordered field containing a predicate defining the integers, but we could include other operations, like exponentials, trigonometric functions, and so on. (We assume, however, that the language is countable, since we are concerned with theories which will ultimately be used to describe physical situations.) This means that the hyperreals closely resemble the usual reals: they form an ordered field, and any operations we have included also extend to the hyperreals.

On the other hand, the hyperreals extend the reals with infinite numbers: there is a real  $\omega$  such that  $r < \omega$  for every actual real number  $r$ . A field with such an element is said to be “non-Archimedean”.

The existence of such systems is promised by the Compactness Theorem for first-order logic. We may take the structure consisting of the real numbers together with addition, multiplication, and any other operations we expect to use, and take the theory  $T$  of all sentences true in this structure. We then add a new constant symbol  $\omega$  and the infinitely many new axioms  $n < \omega$  for each natural number  $n$ . Any finite subset of these new axioms is consistent — to satisfy finitely many axioms, we could interpret  $\omega$  as a big enough natural number. By the Compactness Theorem, it follows that the entire set of new axioms is consistent as well. Any model of the theory of the reals together with these new axioms is then a model of the hyperreals, the constant  $\omega$  must be interpreted as an infinite number, and  $1/\omega$  must be an infinitesimal.

One specific construction of the hyperreals is particularly important. The *ultrapower* construction gives something closer to an explicit construction of the hyperreals. This construction requires the use of an ultrafilter: an ultrafilter on  $\mathbb{N}$  is a collection  $\mathcal{U}$  of subsets of  $\mathbb{N}$ , which we think of as being the “large” sets of natural numbers, such that:

- $\mathbb{N} \in \mathcal{U}$ ,
- $\emptyset \notin \mathcal{U}$ ,
- if  $A \in \mathcal{U}$  and  $A \subseteq B$  then  $B \in \mathcal{U}$ ,
- if  $A \in \mathcal{U}$  and  $B \in \mathcal{U}$  then  $A \cap B \in \mathcal{U}$ ,
- for every set  $A$ , either  $A \in \mathcal{U}$  or  $(\mathbb{N} \setminus A) \in \mathcal{U}$ .

The first four properties are what we expect of a notion of largeness like “with probability 1”: the set of all possibilities has probability 1, the empty set does not, adding events to a probability 1 event still has probability 1, and the intersection of two probability 1 events is also probability 1. The last property is then like requiring that every event have either probability 1 or probability 0.

The principal ultrafilters are the collections  $\{A \mid n \in A\}$  for some  $n \in \mathbb{N}$ . All other ultrafilters are *non-principal* (sometimes called “free”). The existence of non-principal ultrafilters is a consequence of the Axiom of Choice, but cannot be proven in ZF.

An ultrapower of the reals is then a “reduced product” of the reals by an ultrafilter. The key features are that an ultrapower is explicitly described relative to the (highly non-explicit) ultrafilter used to build it, and has an important further compactness property called  $\aleph_1$ -saturation.  $\aleph_1$ -saturation says that whenever we have a countable set of properties  $\phi_n(x)$  so that, for every natural number  $m$ , there is an  $r_m$  with  $\bigwedge_{n \leq m} \phi_n(r_m)$  holding, then actually there must be a single  $r$  so that  $\bigwedge_n \phi_n(r)$  holds. This implies, for instance, that if  $s_1 < s_2 < \dots < t_2 < t_1$  is an increasing sequence of hyperreals followed by a decreasing sequence then there must be a  $u$  in between:  $s_i < u < t_i$  for all  $i$ . (Many such sequences exist in the hyperreals; for instance, consider  $1 < 2 < \dots < \omega - 2 < \omega - 1 < \omega$ .)

A general structure constructed using the compactness theorem need not have this additional property, although ultrapowers are not the only way to obtain  $\aleph_1$ -saturated structures. Nonetheless, ultrapowers are sufficiently better behaved that the term “the hyperreals” is sometimes taken to apply only these structures. (Indeed, this is the context in which the term was coined, [Hewitt, 1948].)

## 2.1 Non-hyperreal infinitesimals

For our purposes, we will consider any non-Archimedean ordered field elementarily equivalent to the reals to be (an example of) the “hyperreals”. However there are non-Archimedean fields which are not hyperreals.

Consider the set of formal fractions of polynomials (with non-zero denominator) in  $\omega$  with real coefficients. Say that two such formal fractions of polynomials are equal iff they are the same when common factors from the numerator and denominator are canceled. One can define addition and multiplication for these fractions of polynomials in the obvious way. Since any polynomial has only finitely many zeros, we can say that a fraction of polynomials is “positive” if its numerator and denominator have the same sign when  $\omega$  is interpreted as a real number greater than any zero of either. Using this, we can say that one fraction of polynomials is greater than another iff their difference is positive, and one can check that with these definitions, this set of formal fractions forms an ordered field, which contains a copy of the reals.

However, this ordered field is not elementarily equivalent to the standard field of reals. For instance, the standard reals satisfy the formula  $\forall x((x > 0) \rightarrow \exists y(y \cdot y = x))$ , stating that every positive element has a square root. However, the formal polynomial  $\omega$  is positive but does not have a square root in this field of formal fractions of polynomials.

For some applications, all that is needed is a non-Archimedean extension of the reals, while for other purposes it matters that the extension be hyperreal. The focus in this paper is on these hyperreal extensions.

## 3 Against applicabilism

### 3.1 What sort of applicability?

Much mathematics is applied in the sciences in ways that are not meant to be taken literally. As Penelope Maddy says about physics, “its pages are littered with applications of mathematics that are expressly understood not to be literally true: e.g., the analysis of water waves by assuming the water to be infinitely deep or the treatment of matter as continuous in fluid dynamics or the representation of energy as a continuously varying quantity.” [Maddy, 1992, p. 281] Similarly, biologists often represent populations of predator-prey systems with the Lotka-Volterra differential equations, using real numbers (rather than integers) to count populations. These models are not taken to accurately represent all the features of waves and populations and the like, but are instead taken to be useful approximations that are much easier to work with. We don’t deny that hyperreals could figure in models of this sort — we merely assert that when they are so used, we should recognize them as idealizations that don’t correspond to the world in the way that other parts of the models do.

The distinction that we aim to make here presupposes some form of scientific realism. An anti-realist who claims that *all* scientific theories are merely models of this sort, with no clear distinction between the representational and the fictional parts of the theories, may deny the cogency of the distinction we are interested in. But if one accepts a distinction between the realism of counting mammals with integers and the instrumentalism of counting mammals with real numbers, then one accepts a distinction of the form we would like to use. Our claim is that *all* uses of the hyperreals in the sciences are at best like the latter.

[Bascelli et al., 2014] motivate the use of hyperreals in science as follows:

there exist physical quantities that are not directly observable. Theoretical proxies for unobservable physical quantities typically depend on the chosen mathematical model. And, not surprisingly, there are mathematical models of physical phenomena which operate with the hyperreals, in which physical quantities take hyperreal values. (p. 850)

The first example of such a phenomenon that they consider is that of Brownian motion, which can be constructed using real or hyperreal numbers. However, we think that *both* of these constructions should likely be taken as idealizations of the physical situation — the relevant motion results from impacts by particles of finite size, with a drop off at sizes below that of a molecule, while either construction of Brownian motion considers all sizes equally. *Either* treatment is better read in an anti-realist way, where it makes useful predictions within a certain range, but doesn’t directly correspond to the world.

We could go through the other examples they discuss, and say which are instances where both the real-valued and hyperreal-valued should be considered in an anti-realist ways, and for which only the hyperreal-valued model has this feature. But in general, we think there are some features that mathematical

models of physical phenomena need to have to be taken in a realist way, and there are principled reasons to think that hyperreal models will usually lack these features.

### 3.2 Automorphisms in applicable mathematics

Suppose that a scientific model has a non-trivial automorphism  $\sigma$ . Then two scientists could apply this model differently: one scientist could predict that some actual measurement will give value  $\alpha$ , while some other scientist could believe that the same measurement should give value  $\sigma(\alpha) \neq \alpha$ . It is clear that neither scientist can be more right than the other: their models are isomorphic. There can be no scientific way to distinguish between a situation in which a measured quantity has the value  $\alpha$  and one where it has  $\sigma(\alpha)$ .

One might still distinguish between these situations by convention: scientific models may require fixing units. For instance, we cannot quite say that lengths are correctly modeled by real numbers without additional information: one scientist could believe that a particular object has length 1, while another believes it has length 3.28084, because the first measures in meters and the second in feet. When we model lengths by the real numbers, we mean this *after* fixing a unit—once we know that a particular object has length 1, we can now uniquely identify all lengths with real numbers.

So an applicable scientific model should be rigid — have no non-trivial automorphisms — once we have fixed a relatively small number of conventional features, like units or coordinate axes. This is essentially the notion of “measurement theory”, as described by [Krantz et al., 1971]. Any quantity that will be represented numerically is said to work in the following way. Some physically meaningful relations and operations are posited, and some empirical generalizations about their behavior are made. A conventional choice is made as to the mathematical relations and operations that will be used to represent the physical ones. Then a theorem is proved showing that if the empirical generalizations hold, then the numerical representation of the physical world is unique up to some particular sort of isomorphism.

For the example of the numerical representation of distance, the physical relation is “longer than” and the operation is the “concatenation” of two segments. We conventionally choose  $>$  as the numerical relation on representations of distances that will correspond to “longer than” for actual distances and  $+$  as the operation on representations of distances that will represent “concatenation” of distances. Given that “longer than” and “concatenation” can be applied to any pair of distances, that every distance can be extended and divided, that “concatenation” and “longer than” commute appropriately, and the Archimedean principle, one can show that a numerical representation of the conventional sort chosen must in fact exist, and this representation is unique up to a choice of which distance corresponds to the length “1”. This is familiar from the fact that different conventions about how to represent distance differ only by a choice of unit (whether the inch or the kilometer or the parsec or the light year). In this sense, distance is said to be on a “ratio scale”.

Similarly, while potential energy is often described as a real valued quantity, it is on a “cardinal scale” (mathematically speaking, it is an  $\mathbb{R}$ -torsor) — only the *differences* between potential energies are meaningful, not the quantities themselves, even after a choice of units. Often a description of a physical situation describes one state as 0 potential energy and quantifies other states relative to that ground state. But two different descriptions might choose two different states to be the ground state. For instance, in describing a satellite orbiting the Earth, one might describe the energy of the orbit with respect to a ground state in which the satellite is on the surface of the Earth, or a ground state in which the satellite is at infinity, or a ground state in which the satellite is in a geostationary orbit. In each case, the assignment of  $\mathbb{R}$ -values as the potential energy of states is a calculational convenience, and the underlying physical property is valued in the coarser structure.

We usually find it most convenient to require these arbitrary choices of units (and a 0 state, for cardinal quantities), so that we can work with the real numbers rather than the more strictly isomorphic mathematical structure. This is even more explicit for higher-dimensional structures. It is often convenient to describe situations in terms of  $x$ - $y$  coordinates, by choosing an origin and axes. But we do not mistake these coordinates for descriptions of reality: they are a computational convenience. In particular, we expect that anything that can be calculated using coordinate systems can also be calculated (perhaps with more difficulty) without them.

Mathematical models with non-trivial automorphisms — like real valued potential energy or coordinate geometry — can be very useful for calculations, but the excess information cannot be physically meaningful. If  $\sigma$  is a non-trivial automorphism and  $\alpha \neq \sigma(\alpha)$ , it cannot be meaningful to say that some measurable quantity has the value  $\alpha$  rather than  $\sigma(\alpha)$ : an isomorphic situation could instead assign it the value  $\sigma(\alpha)$ . The “true” content of the model must be contained in some kind of quotient of the structure, where  $\alpha$  and  $\sigma(\alpha)$  are identified.

This principle gives us a useful tool for interpreting results in computational models with extra automorphisms: we can distinguish between “meaningful” statements — those statements which are automorphism invariant — and meaningless intermediate steps which are not automorphism invariant. For instance, the only statements involving coordinates which we expect to be physically meaningful are those which don’t actually depend on the choice of coordinate system — a physical result which appeared to identify an object’s distance from the  $x$ -axis as an intrinsic fact about the object can be immediately identified as some sort of error.

### 3.3 Applicable structures inside larger structures

Given this stricter notion of applicability, where idealizations don’t count, why should we take the standard reals to be applicable? We have noted that the standard real numbers don’t apply realistically to counting individuals, so perhaps the natural response is to say that *all* uses of the real numbers are similar

idealizations.

[Bascelli et al., 2014] claim that “all physical quantities can be entirely parameterized by the usual rational numbers alone, due to the intrinsic limits of our capability to measure physical quantities.” (p. 853) However, making good on this parameterization using only rational numbers would require strict constraints on what form a physical theory could take. Rational numbers are good enough to report all particular observations to the level of accuracy that we are able to make them. But physical theory goes far beyond particular observations.

As the Pythagoreans discovered already in antiquity, if it is possible for right angles to exist, then there will be very many distances that can not be expressed in terms of integer ratios of some fundamental unit. If physical geometry has a Euclidean region that allows for perpendicular segments of equal length, then there must exist points whose distances have ratio equal to  $\sqrt{2}$ , which is irrational. Similarly, if probability is one of the physical quantities, and it is possible for there to be two independent events that are equally likely, and whose conjunction has probability  $1/2$ , then these events must have probability  $1/\sqrt{2}$ .

These examples so far only motivate the use of quadratic extensions of the rationals, but adequate theorizing about physical laws motivates the use of more complete subfields of the real numbers. For instance, although a two-body system in Newtonian gravitation can have the positions of its elements given at all times by quadratic functions, a three-body system has no closed-form solution for the positions of its elements. Thus, if Newtonian gravitation were a plausible physical theory, we would need to postulate the possibility of distances and times whose relationship to each other is non-algebraic. Or to put it conversely, a physical system that used only algebraic numbers could not be phrased in such an elegant way as the differential equations traditional for Newtonian and later theories.

It may be that the full set of real numbers is unnecessary — perhaps the computable real numbers [Pour-El and Richards, 1989] suffice. However using a superstructure causes no harm to the applicability of the theory so long as the values we attempt to assign to observable quantities belong to the substructure. For instance, even if some field between, say, the algebraic numbers and the computable reals is the correct model of length, we can work in the full model of the reals even though, when assigning values to actual real world quantities, we only ever use values from the subfield.

Because the model is rigid, the extraneous values are still meaningful — the theory still explains what it means for a physical object to have a particular non-computable real as its length, even if this never occurs. Indeed, even if the computable real numbers (or some other substructure of the reals) appear to suffice for all current physics, the discovery of some means of hypercomputation could change the correct choice of substructure.

That is, when an applicable model is a *substructure* of a rigid larger mode, that larger model can remain applicable: at worst it includes theoretical states that can never appear in reality, and which are therefore never assigned to actual quantities.

This is different from the case where an applicable model is a *quotient* of the theoretical model: the latter situation takes the actual physical states and splits them into multiple abstract states. This cannot be (fully) applicable, because every value represents a mix of meaningful information (the choice of an equivalence class) and meaningless information (the choice of an element from an equivalence class). In this case, the model becomes at least partially non-applicable, because every attempt to assign values from the model to observations involves assigning purely abstract information without intrinsic meaning: the only way to obtain meaningful assignments of values is to assign values from the quotient model.

### 3.4 The hyperreals are not applicable

To justify the claim that the hyperreals are an actual representation of some part of the world (rather than merely a useful computational tool), the representation must be unique, or at most involve a small number of arbitrary choices. In a non-hyperreal non-Archimedean structure, it is conceivable that there are multiple units which must be fixed. (Consider an economic model in which there are two “tiers” of desires, with desires in the higher tier lexicographically exceeding any desire in the lower tier; such a model might need two units, one in each tier, to fix a structure up to isomorphism.)

However, if the continuum hypothesis holds, the hyperreals (viewed as an ordered field) have many automorphisms which fix the reals. This continues to hold in any *countably infinite* expansion of the language of ordered fields — for instance, we might want to add many functions arising as solutions to various differential equations, like exponential and trigonometric functions, in addition to countably many units and coordinates. (In an *uncountable* language, the situation is more complicated — indeed, in a large enough language, the hyperreals can become rigid [Enayat, 2006] — but such an infeasible language is, to say the least, an unusual setting for a scientific theory.)

In the absence of the continuum hypothesis, the situation is a bit more complicated: in the absence of the continuum hypothesis, it *is* possible to have rigid structures which could plausibly be called the hyperreals. On the other hand, this also leads to the presence of different, non-isomorphic versions of the hyperreals. A physical theory using the hyperreals would then have to identify a particular one of these structures and explain why that particular structure is the correct one to model the real world in. There is a slim road here through which new information could change our view: new physical discoveries could demonstrate the falseness of the continuum hypothesis and find some way to uniquely distinguish a particular, rigid, hyperreal structure which has an observable physical significance. Because of this possibility, we cannot claim that the hyperreals couldn’t *possibly* be applicable; we can only claim an applicable use of the hyperreals would require substantial new developments in physics.

### 3.5 How close to hyperreals do we need?

Our arguments do not preclude the possibility that some physical quantities are correctly modeled by non-hyperreal non-Archimedean structures extending the reals. To motivate the use of the hyperreals as an actual representation of some part of the world (rather than a useful idealization as suggested above), however, one would need to argue not only that some physical quantity has non-Archimedean behavior, but also that the standard reals are embedded *elementarily* in this structure. Although we are not aware of any physically meaningful quantities that have been argued to have non-Archimedean structure, it doesn't seem out-of-the-question that some quantity could. Some economists and ethicists have argued that there are desires, or moral values, that are lexicographically greater or lesser than others, and the most natural way to represent these involves a non-Archimedean quantity. We could imagine an argument that mass or distance applies in a similarly non-Archimedean way to certain subatomic particles.

But this sort of non-Archimedean structure doesn't require anything as mathematically complex as the hyperreals. As mentioned above, the field of rational functions in one variable over the reals provides a perfectly constructive non-Archimedean structure. To get a hyperreal structure, we must have some argument that the structure is elementarily equivalent to the standard reals despite its non-Archimedean features. If multiplication and addition are the only operations involved, then existence of all square roots and solutions to all odd degree polynomials are sufficient, but most physical quantities appear to be subject to differential equations of some sort as well, which means that elementary equivalence is much stronger.

One might wish to take the hyperreals as a harmless superstructure of the reals, in the same way the reals are harmless superstructure of the computable reals. However the failure of rigidity under the continuum hypothesis means that there is no general explanation of what it can mean for a physical quantity to be appropriately measured by these hyperreals — unlike the rigid case, we do not have an explanation of what it would mean to assign one of these values to an observation. So, as noted above, we believe this can only be done in very narrow circumstances, involving the failure of the continuum hypothesis and an argument picking out a particular rigid version of the hyperreals.

Furthermore, this does not cover the actual proposals which have been made for the uses of hyperreals in the sciences. The uses of the hyperreals generally proposed view the reals as a quotient (technically, subquotient—a quotient of the bounded hyperreals), rather than a substructure. At a model theoretic level, the reals are both a substructure and a certain kind of quotient of a subset. These perspectives lead to different interpretations, however.

A theory which regards the reals as a substructure would take the view that the objects so far observed all have real values—that is, the exact value, as a hyperreal, is one of the hyperreals which happens to have no infinitesimal part; in this interpretation, infinitesimal values are available, but are not needed for any assignments in actual applications.

A theory which views the reals as a (sub)quotient, on the other hand, takes the view that when an observation has a real value, this underspecifies the value: “this object has length 2” means “the length of this object is infinitesimally close to 2”, leaving open which of the many values infinitesimally close to 2 it actually is. This is the perspective taken by actual proposals for the use of the hyperreals in science. It is the more natural use, but we see no way it can be made applicable: in such a theory, every attempt to assign a value to an observation requires picking both a real part of the value and an infinitesimal part, where only the real part (and perhaps some, but not all, of the information in the infinitesimal part) is actually meaningful.

An alternate pathway towards the applicability of mathematics that depends essentially on the Axiom of Choice could involve something like an empirically observable well-ordering of the reals. We think such a possibility is quite implausible, but see no reason to think it should be impossible. One couldn’t directly observe the whole ordering in any finite time, to confirm that it is in fact a well-ordering, but one could hypothetically gather evidence suggesting that it exists. However, such a well-ordering would be quite unlike the sorts of things that we are used to observing.

### 3.6 The recent debate

In recent years, there has been much debate about the value of the hyperreals, with two main views exemplified by Alain Connes and Mikhail Katz, and various coauthors of each. We find much to agree with in what each of these groups say, but we think that each misses some of the points of the others. We hope that by framing the discussion as concerning two distinct types of realism, we can keep what is valuable from each side.

[Kanovei et al., 2013, p. 261] say “Connes’ variety of Platonism can be characterized more specifically as a prescriptive Platonism, whereby one not merely postulates the existence of abstract objects, but proceeds to assign “hierarchical levels” (see [Connes et al., 2001], p. 31) of realness to them, and to issue value judgments based on the latter.” We don’t want to make value judgments, or say that the higher levels of the hierarchy are less real — we just want to warn people to be careful in how they use them.

For example, models in which probabilities have hyperreal values can provide insight into probabilistic paradoxes, as in [Bartha and Hitchcock, 1999]. (This is discussed further in the last section of this paper.) However the content of such a model cannot really be in claiming that one should truly have hyperreal degree of belief in something. Rather, a model which uses the hyperreals is implicitly making claims about the way and the order in which various limits approach infinity. The hyperreals provide a useful mathematical tool for working with these limits, but any analysis of the suitability of such a model has to include unpacking and examining these choices about limits.

One way this distinction is discussed is in a supposed contrast between mathematical entities that are “mere tools” and ones that are real “objects” of study. We don’t want to make any absolute version of this distinction, but think it is

useful when relativized to a particular context. The complex numbers might be treated as mere tools in some contexts (such as understanding the distribution of prime numbers), but in other contexts (such as thinking about algebraically closed fields), complex numbers are the most natural object of study. Conversely, while the hyperreals are meaningful objects worthy of their own study, there are other contexts (particularly dealing with real analysis, and the physical theories using it) where they are mere tools for avoiding some complexity in dealing with quantifier structure.

### 3.7 Can hyperreals be “named”?

One major line of debate focuses on the claim that no non-standard hyperreal can be “named”. Connes tries to establish this claim by arguing that a given non-standard integer in a hyperreal field can be used to generate a non-principal ultrafilter over the natural numbers, or a non-measurable set of reals. Since these complex entities seem to be beyond some limit of complexity for physical beings like us to grasp, this is said to raise problems for any appeal to hyperreals.

Kanovei, Katz, and Morman dispute Connes’ claims about the association of a non-principal ultrafilter with a given hyperreal field. (pp. 272-275) They point out that some of his claims are ambiguous, depending on whether we consider the hyperreal to be given as an isomorphism type (in which case they argue that it is false) or to be given as an explicit ultrapower of the standard reals (in which case it is trivial to be able to extract an ultrafilter from it). Furthermore, they point out that in  $L$ , there is a definable non-measurable set (p. 278) so that the use of a hyperreal to define a non-measurable set doesn’t rule out the definability of a particular hyperreal. They appeal further to a result of [Kanovei and Shelah, 2004] showing that in fact, ZFC entails the existence of a definable hyperreal field.

We claim that definability is not the real issue here. These things are “definable” in the strict mathematical sense, but the “definitions” don’t serve the purpose we need definitions to serve in physical models, of making it possible to uniquely measure quantities. The Kanovei and Shelah construction relies on an ultrafilter over a set that itself has no definable elements! Even if it is possible to define a particular ultrafilter, and use it to define a hyperreal field, and then define a particular non-standard element of this field, these definitions don’t appeal to the distinctions that we care about in application. This shows up in the fact that there are automorphisms preserving all the features of the model that are relevant in application, that nevertheless move any given non-standard element. Thus, the features that are used in the definition are not relevant in application.

## 4 In Defense of Factualism

We have no such concerns when using hyperreals and the Axiom of Choice within mathematics. Indeed, we think the validity of the Axiom of Choice and

the hyperreals as a tool in mathematics is unavoidable.

We don't have much to add to existing philosophical defenses of the Axiom of Choice and its consequences. We don't expect the arguments we give to convince any opponents of the factualism of the Axiom of Choice. But we want to make clear that we do endorse it.

Some defenses of the Axiom of Choice focus on the fact that (in at least some forms) it seems intrinsically plausible — for instance, in the form that the product of a non-empty collection of non-empty sets is non-empty. Other defenses point out the extrinsic virtues it has of unifying results from various areas of mathematics and systematizing them in various ways. (Both of these types of justification are discussed in [Maddy, 1988, pp. 487-9].) Platonists about mathematical entities emphasize that these methods of reasoning are part of our reasoning about the ordinary physical world, and claim that there is no principled reason why these methods of reasoning shouldn't be useful in the mathematical realm. But fictionalists and others that deny the literal truth of mathematics also often accept that there is a correct natural way to develop the story of mathematics, and accept the same theoretical virtues that motivate the platonist.

One major point we want to make in this paper is that some reasons that people have given for rejecting the Axiom of Choice don't really require that we reject it — they just require some care in deciding which applications of mathematics to the physical world to take literally, as discussed above. But another point is that even someone who denies the factual truth of the Axiom of Choice should be willing to accept many of the applications of it that we make. While we accept the factual truth of the Axiom of Choice, the rest of this section is devoted to discussion of some uses that don't even require that.

A variety of results [Henson and Keisler, 1986, van den Berg et al., 2012, Sanders, 2016, Nelson, 1988, van den Berg and van Slooten, 2017, Goodman, 1978, Sanders, 2015] show that, in many settings, these tools are *conservative*: adding them to a formal system does not give new arithmetic consequences. Moreover, many of these proofs are themselves quite concrete, including several giving explicit, constructive translations which convert proofs which use the hyperreals to proofs which do not.

In the presence of these results, the dispute between the realist mathematician who uses the hyperreals to prove a fact about the natural numbers and the anti-realist mathematician, who considers the Axiom of Choice meaningless, is a dispute over terminology, not mathematics.

For example, the nonstandard analyst would say that a function  $f : \mathbb{R} \rightarrow \mathbb{R}$  is continuous at  $x \in \mathbb{R}$  if for every  $y$  such that  $x - y$  is infinitesimal (written  $x \simeq y$ ),  $f(x) - f(y)$  is also infinitesimal. In quantifiers, this says  $f$  is continuous at  $x$  if

$$\forall y \in \mathbb{R}^*[x \simeq y \rightarrow f(x) \simeq f(y)].$$

However there is a systematic, syntactic way to extract the standard content out of this statement. We observe that  $x \simeq y$  is the same as saying that for

every positive rational  $\delta > 0$ ,  $|x - y| < \delta$ , so saying  $f$  is continuous at  $x$  is saying

$$\forall y \in \mathbb{R}^* [(\forall \delta \in \mathbb{Q}^+ |x - y| < \delta) \rightarrow (\forall \epsilon \in \mathbb{Q}^+ |f(x) - f(y)| < \epsilon)].$$

We can then use ordinary quantifier manipulation on first-order sentences to see that this is equivalent to:

$$\forall y \in \mathbb{R}^* \forall \epsilon \in \mathbb{Q}^+ \exists \delta \in \mathbb{Q}^+ [|x - y| < \delta \rightarrow |f(x) - f(y)| < \epsilon].$$

Finally, the saturation of the hyperreals allows us to turn initial nonstandard quantifier into a standard quantifier while pushing it inside—this statement is equivalent to:

$$\forall \epsilon \in \mathbb{Q}^+ \exists \delta \in \mathbb{Q}^+ \forall y \in \mathbb{R} [|x - y| < \delta \rightarrow |f(x) - f(y)| < \epsilon].$$

That is, the nonstandard analyst and the conventional mathematician outright entirely agree on the definition of continuity, they merely express the notion in different notation. (Sanders expands on this point in [Sanders, 2016].)

This same method—with additional complications—applies more generally, including to results originally formulated in the nonstandard language. Consider an example from the mathematical literature: Jin’s nonstandard proofs of some results about density in sets of natural numbers [Jin, 2001]. Jin considers subsets of the natural numbers, and, several notions of density, including *Banach density*, given by

$$BD(A) = \lim_{n \rightarrow \infty} \sup_{0 \leq k \leq m, m-k=n} \frac{|A \cap [k, m]|}{n+1}$$

and *lower asymptotic density*

$$\underline{d}(A) = \liminf_{n \rightarrow \infty} \frac{|A \cap [0, n]|}{n}.$$

Jin’s key lemma, in the terminology of nonstandard analysis, says

Suppose  $A \subseteq \mathbb{N}$  and  $BD(A) = \alpha$ . Then there exists an interval of hyperfinite length  $[H, K]$  such that for almost all  $x \in [H, K]$ ,  $\underline{d}(*A - x) = \alpha$ .

Even to the mathematician who considers the existence of the hyperreals to be meaningless, this is still a perfectly reasonable mathematical statement: it is equivalent to saying

Suppose  $A \subseteq \mathbb{N}$  and  $BD(A) = \alpha$ . Then for any  $g$  and any functionals  $E, D$  and  $N$ , there exists functions  $M, H, K$  such that, taking  $\epsilon = 1/E(M, H, K)$ ,  $\delta = 1/D(M, H, K)$ , and  $n = N(M, H, K)$ ,  $h = H(\epsilon, \delta, n)$  and  $k = K(\epsilon, \delta, n)$ , we have  $k - h = g$  and for all but  $\epsilon(k - h)$  elements  $x \in [h, k]$ , there is an  $m \in [n, M(\epsilon, \delta, n)]$  such that  $\frac{|(A-x) \cap [0, m]|}{m} \geq \alpha - \delta$ .

In this latter statement, all functions are computable and all existence statements fully constructive, so there should be no disagreement about its meaningfulness. The anti-realist may find Jin’s formulation infelicitous, given its apparent reliance on the existence of hyperreals, but this reliance is only apparent—when Jin’s version says “there exists an interval”, the anti-realist must understand this as meaning something other than existence in a conventional sense, but this is only a difference of terminology.

Here we agree with Bishop, who remarked that Keisler’s calculus book using the hyperreals “offers no evidence that the hyperreal numbers are anything except a device for proving theorems about the real numbers” [Bishop, 1977]. Connes makes a similar point. But unlike Bishop and Connes, we do not see this as a criticism — when one’s goal is to prove theorems, devices for doing so are all to the good. Indeed, this is precisely the strength of the hyperreals: they are a linguistic tool for keeping track of “epsilonics”, and quantifiers more generally, in a cognitively easier way without fundamentally changing our understanding of the reals.

Some applications of the hyperreals to probability have this form. For instance, Bartha and Hitchcock deal with the troublesome “shooting-room paradox” in probability by using the hyperreals to create a standard real-valued probability function that fails to be countably additive, but has other nice features that resolve the paradox.

We stress that while we will be using non-standard analysis as a tool, the probability measures that we ultimately define will be strictly real-valued, and finitely additive. Thus we are not committed to the existence of infinitesimal degrees of belief or anything of that sort. Just as imaginary numbers can be used to facilitate the proving of theorems that exclusively concern real numbers, our use of nonstandard analysis will be used to facilitate and motivate the construction of purely real-valued measures. [Bartha and Hitchcock, 1999, p. 416]

E.T. Jaynes, in Chapter 15 of his book, [Jaynes, 2003] argues that (at least in probability theory) one should always work only with finite sets and consider infinite versions of situations only as limits of finite versions of them. He argues that violations of conglomerability or countable additivity only arise when one disobeys these suggestions and tries to directly put a probability measure on a set already assumed to be infinite. Various paradoxes of probability involve considering an infinite probability space that is a limiting version of multiple different types of finite probability spaces, and tries to apply intuitions from each in a conflicting way.

Jaynes’s limiting advice seems like a non-starter for any modern mathematician who is untroubled by the notion of infinite sets as definite, well-behaved objects. Thus, we don’t agree with his advice that *probability* theory should never be applied to them, except perhaps by a special terminological usage that the word “probability” be restricted to a certain type of reasoning about physical situations. And we emphatically disagree with his dismissive approach to much mainstream statistical practice:

Thus it is not surprising that those who persist in trying to evaluate probabilities directly on infinite sets have been able to study finite additivity and nonconglomerability for decades — and write dozens of papers of impressive scholarly appearance about it. Likewise, those who persist in trying to calculate probabilities conditional on propositions of probability zero, have before them an unlimited field of opportunities for scholarly looking research and publication — without hope of any meaningful or useful results. (p. 1528)

His viewpoint is akin to that of constructivists that deny the meaningfulness of the Axiom of Choice because of the essential role infinite sets play in interpreting it.

However, in generalized mathematical probability theory, we think that Jaynes’s suggestion of keeping track of the limiting operations is an important one. Bartha and Hitchcock show simply that any countably-additive distribution will violate one of the intuitions of the shooting-room paradox, even though it gives an internally consistent set of results (as Jaynes suggested). However, their use of the hyperreals allows them to take limits in a different order, and keep all the intuitions of the problem, while ending up with a probability distribution that is only finitely (rather than countably) additive. They also show that their distribution could be generated using only standard real numbers, by keeping very careful track of limits, though the use of hyperreals makes it easier.

It is crucial to the distinction we wish to draw between the mathematical power of the hyperreals and their physical applicability to notice that while statements using the hyperreals are meaningful, the meaning is something that could be expressed without them. We can compare this to the value of a coordinate system in geometry: the choice of a non-canonical coordinate system is a calculation tool, and is only useful if what comes out at the end is independent of the coordinate system. Similarly, working with the hyperreals involves highly non-canonical choices, and is meaningful if, but only if, what comes out at the end is independent of those choices. And as we saw in the previous section, attempts to apply hyperreals directly as representatives of particular states of the world (rather than as tools for generating a real-valued representation) will generally not be independent of the non-canonical choices.

## References

- [Bartha and Hitchcock, 1999] Bartha, P. and Hitchcock, C. (1999). The shooting-room paradox and conditionalizing on measurably challenged sets. *Synthese*, 118:403–437.
- [Bascelli et al., 2014] Bascelli, T., Bottazzi, E., Herzberg, F., Kanovei, V., Katz, K. U., Katz, M. G., Nowik, T., Sherry, D., and Shnider, S. (2014). Fermat, Leibniz, Euler, and the gang: The true history of the concepts of limit and shadow. *Notices of the American Mathematical Society*, 61(8):848–864.

- [Bishop, 1977] Bishop, E. (1977). Review: H. Jerome Keisler, Elementary calculus. *Bull. Amer. Math. Soc.*, 83(2):205–208.
- [Chakravartty, 2017] Chakravartty, A. (2017). Scientific realism. *Stanford Encyclopedia of Philosophy*.
- [Connes et al., 2001] Connes, A., Lichnerowicz, A., and M., S. (2001). *Triangle of Thoughts*. American Mathematical Society.
- [Enayat, 2006] Enayat, A. (2006). Automorphisms of nonstandard reals, revisited. Foundations of Mathematics mailing list.
- [Goodman, 1978] Goodman, N. D. (1978). Relativized realizability in intuitionistic arithmetic of all finite types. *J. Symbolic Logic*, 43(1):23–44.
- [Henson and Keisler, 1986] Henson, C. W. and Keisler, H. J. (1986). On the strength of nonstandard analysis. *J. Symbolic Logic*, 51(2):377–386.
- [Hewitt, 1948] Hewitt, E. (1948). Rings of real-valued continuous functions. I. *Trans. Amer. Math. Soc.*, 64:45–99.
- [Jaynes, 2003] Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press.
- [Jin, 2001] Jin, R. (2001). Nonstandard methods for upper Banach density problems. *J. Number Theory*, 91(1):20–38.
- [Kanovei et al., 2013] Kanovei, V., Katz, M. G., and Mormann, T. (2013). Tools, objects, and chimeras: Connes on the role of hyperreals in mathematics. *Foundations of Science*, 18:259–296.
- [Kanovei and Shelah, 2004] Kanovei, V. and Shelah, S. (2004). A definable nonstandard model of the reals. *Journal of Symbolic Logic*, pages 159–164.
- [Krantz et al., 1971] Krantz, D., Luce, D., Suppes, P., and Tversky, A. (1971). *Foundations of Measurement*, volume 1. New York Academic Press.
- [Linnebo, 2013] Linnebo, O. y. (2013). Platonism in the philosophy of mathematics. *Stanford Encyclopedia of Philosophy*.
- [Maddy, 1988] Maddy, P. (1988). Believing the axioms. i. *The Journal of Symbolic Logic*, 53(2):481–511.
- [Maddy, 1992] Maddy, P. (1992). Indispensability and practice. *The Journal of Philosophy*, 89(6):275–289.
- [Nelson, 1988] Nelson, E. (1988). The syntax of nonstandard analysis. *Ann. Pure Appl. Logic*, 38(2):123–134.
- [Pour-El and Richards, 1989] Pour-El, M. B. and Richards, J. I. (1989). *Computability in analysis and physics*. Perspectives in Mathematical Logic. Springer-Verlag, Berlin.

- [Robinson, 1996] Robinson, A. (1996). *Non-standard analysis*. Princeton University Press.
- [Sanders, 2015] Sanders, S. (2015). The unreasonable effectiveness of nonstandard analysis.
- [Sanders, 2016] Sanders, S. (2016). The computational content of nonstandard analysis. In *Proceedings Sixth International Workshop on Classical Logic and Computation*, volume 213 of *Electron. Proc. Theor. Comput. Sci. (EPTCS)*, pages 24–40. EPTCS, [place of publication not identified].
- [van den Berg et al., 2012] van den Berg, B., Briseid, E., and Safarik, P. (2012). A functional interpretation for nonstandard arithmetic. *Ann. Pure Appl. Logic*, 163(12):1962–1994.
- [van den Berg and van Slooten, 2017] van den Berg, B. and van Slooten, L. (2017). Arithmetical conservation results. *ArXiv e-prints*.