Problems in the philosophy of mathematics: A view from cognitive science

Steven T. Piantadosi

May 10, 2015

The success of mathematical thinking relies in great part on the creation of a language for expressing precise ideas. This language contains in the simplest cases notions like disjunction, implication, and the concepts of calculus, and in most complex cases the vocabulary of group theory or the entire *Principia Mathematica.* Mathematics must create its own language because natural language is not up to the task—natural language is loaded with ambiguity, vagueness, and conceptual representations that decades of effort in cognitive science have yet to adequately formalize. Even when mathematicians borrow natural language terms, they must refine the meanings in order to leave no uncertainty about what was intended. Mathematical language explicitly distinguishes, for instance, "if" vs. "iff", "or" vs. "xor," and a meaning of "some" that ignores the awkwardness in natural language of using it when "all" is also true.

The primary argument of this paper is that at least a few problems in philosophy—including some in the ontology of mathematics—arise when the precision of natural language is overestimated and we mistake the fuzziness of natural linguistic and conceptual systems for the precision needed for mathematical ones. Attempting to answer natural language questions about the reality of mathematical objects is as futile as trying to decide whether it is really true that "Bill or John played the accordion" when they both did. This question is not about the nature of the universe; it is at best a question about what "or" means to people or how "or" is most naturally used. The problem becomes even worse for the topics in mathematical ontology—questions like existence, sameness, and truth—which sound reasonable in natural language, but which are imprecise at their core. This critical view follows a rich tradition in philosophy, touching on analytic philosophy, ordinary language philosophy, and Logical Positivism in general, as well as the works of Bertrand Russell, Rudolf Carnap, and Ludwig Wittgenstein in particular.

My article will attempt to provide a cognitive perspective on these issues, arguing that there is now substantial evidence that natural human concepts are imprecise, and in fact imprecision is likely useful to their normal usage and in the context in which we evolved. At the same time, we appear to have poor metacognition about the precision of our concepts, meaning that we view our own terms and concepts are not being inherently vague or imprecise. These two facts conspire create what I will call an *illusion of conceptual precision.* The strength of this illusion makes it very easy for us to ask questions that seem answerable—perhaps even deep—but have no substance beyond our own fuzzy mental constructs.

Human concepts are fuzzy, probably for a good reason

Natural language provides convenient labels for concepts and 50 years of cognitive science has yet to make substantial progress on formalizing what type of thing concepts might be. One important debate concerns whether concepts are represented with logical rules, stored examples (called exemplars or prototypes), or something else entirely. Proponents of the rule view call on our ability to represent abstract, logical concepts like *even numbers*, which appear to be most naturally formalized in a logic-like language, not unlike that found in mathematics. Such a *Language of Thought* [12] would explain humans' systematic patterns of thought, as well as the structured, compositional form that many of our beliefs take [13]. For instance, we are able to think about our own thoughts ("I doubt I believed yesterday that Sam would eat a mouse.") and construct richly interconnected representations in natural language ("My doctor's sister's husband never forgets to take a pair of glasses with him when he visits a town with at least one movie theater."). It appears difficult to explain our ability to think such thoughts without hypothesizing a structured representation system for mental language, perhaps one like those used in logic [26]. If mental representations were logical structures, that might give hope that mental concepts could be firmly grounded with the same precision and quality of meaning as mathematical systems.

Unfortunately, cognitive systems appear to be much more complex, particularly below the level of sentences. Although a few words are amenable to rule-like definitions, there are probably not many [12]. Philosopher Jerry Fodor even argues that even a term like the transitive verb "paint" cannot be captured with a definition even if it is taken for granted that the noun "paint" is known [14]. It is tempting to think that "X paints Y" if X covers Y with paint. But, as Fodor points out, this won't work because if a paint factory explodes, it can cover something with paint without painting it. It seems that X has to be an agent. But then, an agent could accidentally cover their shoes with paint, in which case it wouldn't be quite right to say "X painted their shoes." There seems to be a necessary notion of intention—"X paints Y" if X is an agent who covers Y with paint and X intends to do so. But then Fodor considers Michaelangelo dipping his brush in paint. He doesn't seem to have painted *the brush* even though he intended to cover it with paint. Fodor remarks, "I don't know where we go from here."

And even for concepts that seem to have natural definitions, the definitions still do not capture much of humans' knowledge. A classic example is that of a "bachelor," which could naturally be defined as an unmarried male. Critics of the definitional view point out that even rule-like concepts appear to have "fuzzy" boundaries. Most readers would agree that although a widower may technically be a bachelor, he certainly isn't a good one. The pope, another unmarried male, is another bad example. My baby nephew isn't really a bachelor yet, either, and—even stranger—it's not clear at what age he might become one. These examples demonstrate that we know a lot more than "unmarried male"—we know about typical and atypical features, as well as different senses, connotations, and analogies. Strikingly, the fuzziness of our concepts applies even to *mathematical* ones: subjects report, for instance that 2 is a "better" even number than 178 [1][1]. Even experts in mathematics would probably agree that $f(x) = x^2$ is a "better" function than the Dirichlet function.

Motivated by these effects, many theories of conceptual representation do not rely on systems of precise logical rules. Instead, categorization and conceptualization may work based on the similarity to stored examples, as in *prototype theory* (using a single stored example) or *exemplar theory* (using many) (for an overview see [24]). A new example could then be categorized by measuring its similarity to the stored example(s). In this case, we might not have a precisely defined notion of bachelor, but could still decide whether someone was a bachelor by comparing them to previous bachelors we have seen and remembered. In this case, there may be no fact of the matter—only varying degrees of bachelorhood.

Of course, prototype and exemplar theories also have challenges (see [24]). One is in particular with explaining the aspects of thinking which look structured and compositional (e.g., how might prototypes be combined to represent meanings like "short giraffe"?). Another is in explaining how our rich system of concepts interrelate (e.g., are the terms "parent" and "child" defined in terms of each other?). A third concerns how one might explain concepts that really do seem definitional, and whose features (surface similarity) are not diagnostic: for instance, an old woman with gray hair may be very similar to any previously seen grandmothers yet not be one.

The correct theory of conceptual representation must make peace with all sides of human thinking—its rule-like nature, as well as its fuzzy, approximate, and graded nature.[2] A fair generalization from this debate is that many of the concepts we use everyday seem to be imprecise, although the degree and nature of the imprecision varies considerably. Critically, if an imprecise system forms the foundation for cognition, there is little hope for formalizing concepts and determining fundamental truth behind many questions that can only be phrased in natural language. This is not to disparage them—human concepts and language have given rise to remarkable cognitive, cultural, and social phenomena.

The prevalence of conceptual fuzziness is why one could write a dissertation on the definition of a "kill," "treason," or "blame." This graded nature of concepts is reflected in our linguistic constructs, including phrases like "he's a pretty good chef" and "he's mostly trustworthy," and "she's probably your mother." The problem is particularly worrisome for the terminology that forms the foundation of science—we don't really have a precise notion of what makes something a cause[3], a computation, a representation, or a physical entity, despite the centrality of these notions for scientific theories. Difficulties with aligning our intuitive concepts with those that best describe the objective world can also be see in debates about the definition of races, species, life, and intelligence. Imprecision is why it is unclear when a fetus becomes a human, and why there is debate about whether viruses are alive. Even the biological definition of "male" and "female" is surprising—it is not based on the readily apparent

[1]Interestingly, these results were used to argue that typicality effects were poor evidence *against* definitions since graded judgments can appear for these clearly definitional concepts.

[2]Some exciting recent work unifies stochastic conceptual representations with structure, Turing-completeness, and inference [17]

[3]Though see [31] for mathematical work formalizing this notion.

physical differences that our intuitive concepts might use. Instead, the biological definition is based on the size of gametes produced: females in a recently discovered species of insect use a "penis" (apparently another vague concept) to extract sperm from the male [48].

There is no accepted account of why natural language concepts lack precision. It could be that such undefined concepts are actually more useful than precise ones, in the same way that ambiguity permits efficient communication by allowing us to leave out redundant information [34]. Vague terms certainly permit communicative analogies and metaphors ("All politicians are criminals."). Alternatively, fuzziness might result from the necessities of language learning: the average child will acquire about 10 words per day on average from age 1 to 16. Perhaps there is not enough time or data to acquire specific, detailed meanings. Indeed, experts often come to have more refined word meanings in their domain of expertise. For instance, "osmosis" and "diffusion" are used interchangeably by many, except experts in biology or chemistry. This is consistent with the fact that meanings are always changing, showing that cultural transmission of meanings is noisy and imperfect. Despite the efforts of linguistic prescriptivists—the climate-change deniers of the language world—the word "literally" has come to mean "figuratively" just as "you" has come to mean "thou." Such linguistic change is gradual, inevitable, and self-organized, meaning that children can't be acquiring precise replicas of their parents' conceptual and linguistic systems.

Alternatively, human language and concepts may be imprecise due to abstraction. We are able to have thoughts about people, events, and objects at multiple levels: I can think of my family as a family, a collection of individuals that obey certain relations, a collection of undetached-person-parts (Quine), a curious collection of symbiotic cells (Margulis), as DNA-bound information trying to replicate itself (Dawkins), or as extraordinarily complex interactions of atoms. Our thinking admits such varying levels of abstraction because they are necessary for understanding the causal relations in the world. Hilary Putnam's example is of the inability to put a square peg in a round hole: even though this fact about the world results from "just" the interaction of atoms, understanding of *why* this is true comes about from more abstract—and nonreductionist—notions of geometry [37]. Steven Pinker's example is that of World War I which—though it is "just" a complex trajectory of subatomic particles—is only satisfactorily explainable in terms of high-level concepts like nations and treaties.

Once the leap has been made to allow more abstraction, concepts quite naturally come to lose some of their precision. An abstract concept deals with generalities and not particulars. As a result, it seems very natural for systems of abstract representation to break down in unusual particular instances since the particulars are not tracked in detail. Moreover, the abstraction must be discovered through inference and generalization. This means that abstractions may additionally be fuzzy since they represent parts of reality that are not directly experienced.

The illusion of conceptual precision

Even as human cognition is a remarkable force of nature, human *meta*cognition often leaves much to be desired. Metacognition refers to our ability to understand the abilities and limitations of our own cognitive skills. In many domains, we perform poorly on metacognitive tasks, often overestimating our own capacity. 93% of US drivers believe themselves to be above average, for instance [45]. Patricia Cross [7] found that 94% of faculty members rate themselves as above-average teachers, noting that therefore "the outlook for much improvement in teaching seems less than promising." People often believe that they understand even commonplace systems like locks, helicopters, and zippers much better than they do, a phenomenon called the "illusion of explanatory depth" [39]. An "illusion of understanding" in political policies has even been linked to political extremism [11].

Such metacognitive limitations likely extend to the core properties of our own concepts. This is why the marginal cases of "bachelor" are hard to think of—much easier to think of are the more typical bachelors. In many cases, our explicit knowledge of concepts appears to be much more like a knowledge of the average example rather than the decision boundary. This observation is a fairly nontrivial one in cognitive science: many classification techniques (e.g., support vector machines) represent only the boundaries between classes. Human cognition probably does not use such discriminative concepts, as we both don't know where the boundaries are very well. This is exactly what would be predicted under a prototype, exemplar, or generative view—we tend to know category centers or typical cases. This view is supported by the *typicality effect* in word processing, where typical examples tend to receive privileged processing, including faster response times, stronger priming, and most critically a higher likelihood of being generated in naming tasks [38]. Thus, it is very easy to describe typical bachelors or typical birds; and less easy to construct atypical ones. Human conceptualization is focused on the parts of concepts that are categorized well, not the parts that are iffy or near the boundaries. This property may make sense if most of our uses of concepts are in labeling or referring. Perhaps it is not important to know exactly where the boundaries are when producing labels (words), since by the time you get to the boundaries of bachelors (e.g., the pope) there is often a better term available (e.g., "the pope") that we can use instead.

One unfortunate consequence of this illusion is that we often get lost arguing what a term *really* means, instead of realizing that any definition we pick will be arbitrary in some sense. This can be seen in debates about the definition of vulgarity or profanity, for instance. An inability to formalize our intuitive notions was once celebrated in the Supreme Court: in a famous case on the definition of pornography, Justice Potter Stewart famously declared "I shall not today attempt further to define the kinds of material ... and perhaps I could never succeed in intelligibly doing so. But I know it when I see it, and the motion picture involved in this case is not that." The "I'll know it when I see it" doctrine now has its own Wikipedia page and has been heralded as an exemplar of judicial commonsense and realism—even though it fails completely to solve the problem.

One remarkable failure of conceptual metacognition was the recent public uproar over changing the definition of "planet" in a way that excluded Pluto. The decision was simply one about the *scientific* definition of a term, and has no bearing on the public use of the term.[4] However, the public had very strong views that *of course* Pluto was a planet and that should not change. The problem arose because people believed that they had a notion of planethood which was correct (perhaps even objectively) and therefore there was something wrong with the changed definition. People's cognitive representation of a planet was not precise—I'll know a planet when I see one!—and this fact clashed with their strong beliefs about Pluto. The precise (scientific) definition[5] *happened* to demote Pluto from "planet" to "dwarf planet." The only way you could think the definition was "wrong" would be to beg the question, or to believe—almost certainly erroneously—that your cognitive representation was a better scientific definition than the astronomers'.

Curious conceptual metacognition also appears in children's development. Many humans—children included—are *essentialists* meaning that we believe there is some unobservable essence to our psychological categories (see [15]). For instance, a recent buyers were willing to pay over 1 million dollars for JFK's golf clubs, even though there is nothing physically different about the clubs that ties them to JFK. Instead, they seem to have some nonphysical essence linking them to JFK simply because he once possessed them. A case more typical to development is the difference between men and women: children may believe there to be an inherent difference, even though they do not know what it is. In development, such essences may be "placeholders" for things children do not yet understand [25], suggesting that our concepts may not even at their core be about properties of objects [15]. A concept can apparently be formed even in the absence of any real notion of what the concept might be.

Our sense of having precise and justifiable concepts when we really do not should be familiar in both mathematics and philosophy, where we often learn with great pains that we don't fully understand a concept. The concept of a *real number* is a nice example, where students will often get some conception of this term from high school. But a full, deep understanding of its properties requires more thought and study—and it's amazing to learn the number of things we *don't* know about real numbers once we first have the concept. Such non-obvious properties include, for instance, knowledge that decimal expansions of real numbers are not unique ($0.9999\ldots = 1$), knowledge that the cardinality of real numbers is a different infinity than the reals, knowledge that *almost every* real number is uncomputable and incompressible, a randomly sampled real number (from, e.g., a uniform distribution on $[0, 1]$) is certain to have probability 0 of being chosen, and knowledge that any Dedekind-complete ordered field is essentially the same as real numbers. Internet "debates" about these topics make fascinating reading relevant to the illusion. Deep understanding of concepts requires much more than our easily achieved conceptions about the world. But, we should not take for granted that it is possible at all.

The relevance of the illusion to philosophy

The toxic mix of fuzzy concepts and illusory precision is relevant to areas of philosophy that state their problems in natural language. This problem has been recognized by philosophers. In his *Tractatus Logico-Philosophicus*, Wittegenstein writes,

[4]Under common usage, some people—most recently some viral-internet QVC hosts—believe that the moon is a planet. It can be forgiven since the moon doesn't seem like a "satellite" much anymore.

[5]The International Astronomical Union holds that a planet must (i) orbit the Sun, (ii) be nearly round, and (ii) have "cleared the neighborhood" around its orbit.

> 4.002 Man possesses the capacity of constructing languages, in which every sense can be expressed, without having an idea how and what each word means—just as one speaks without knowing how the single sounds are produced. ...
>
> Language disguises the thought; so that from the external form of the clothes one cannot infer the form of the thought they clothe, because the external form of the clothes is constructed with quite another object than to let the form of the body be recognized. ...
>
> 4.003 Most propositions and questions, that have been written about philosophical matters, are not false, but senseless. We cannot, therefore, answer questions of this kind at all, but only state their senselessness. Most questions and propositions of the philosophers result from the fact that we do not understand the logic of our language. ...
>
> And so it is not to be wondered at that the deepest problems are really no problems.
>
> 4.0031 All philosophy is "'Critique of language" ...

Rudolf Carnap's *Pseudoproblems in Philosophy* explains away philosophical problems like realism as simply insufficiently grounded in a scientific epistemology to even have meaning [5]. He writes of such philosophy, "A (pseudo) statement which cannot in principle be supported by an experience, and which therefore does not have any factual content would not express any conceivable state of affairs and therefore would not be a statement, but only a conglomeration of meaningless marks or noises."

Daniel Dennett describes the term "deepity" for phrases which sound profound but are in actuality meaningless [9]. A nice prototype is the sentiment "Love is only a word," which is true on one reading and seems deep on another, until you think about it. Deepak Chopra is a purveyor of household deepities, including "The idea of a world outside of awareness exists in awareness alone," "Existence is not an idea. Only its descriptions are," and "Only that which was never born is deathless." The ability to form such meaningless thoughts that sound profound is one consequence of the illusion of conceptual precision. The cultural celebration of such sentiments makes those inclined towards precision sympathetic with Hunter S. Thompson, whose response to "What is the sound of one hand clapping?" was to slap the questioner in the face [46].

The effect of imperfect meanings in natural language can also be seen in paradoxes of language and logic. One is Bertrand Russell's Barber paradox [40], concerning the barber who cuts a person's hair if they do not cut their own. Does the barber cut his own? In set theory, the paradox is to consider set S that contains all sets that do not contain themselves ($S = \{s \ s.t. \ s \notin s\}$). Does S contain itself? This problem provided a challenge to phrasing set theory in terms of natural language (*naive set theory*) as well as other formal systems. Such logical catastrophes demonstrate that natural language is not precise enough for sound reasoning: it is too easy to construct sentences whose meaning *seems* precise, but which cannot be. Use of underspecified natural language concepts yields other puzzles, including the proof that there is no uninteresting integer. For if there were, there would have to be a smallest uninteresting integer and sure that number would be interesting (since it is the smallest uninteresting integer).[6] There is also the smallest number that cannot be described in less than sixteen words (yet I just described it). Relatedly, there are even sentences—called Escher sentences (see [36])—which give the illusion of having a meaning, when really they don't: "More people have eaten tomatoes than I have." This is a semantic illusion in the same way that a Necker cube is a visual illusion and Shepard tones are an auditory illusion.

Not only are many of the terms of debate imprecisely grounded, but the illusion creates a social or cognitive pull that makes humans particularly willing to engage in meaningless discussions: we do not easily recognize our own imprecision. Here, I'll consider three examples: the existence, sameness, and nature of mathematical concepts. In each case, I'll argue that because the underlying language is natural language, core questions in these domains have yet to be usefully posed. These questions are as useful to argue about as whether Russell's barber *can* or *cannot* cut his own hair.

Existence

One of the core questions of mathematical ontology concerns the question of whether mathematical objects like numbers exist. The range of philosophical positions that have been argued for range from

[6] In 2009, Nathaniel Johnston found that 11630 was the smallest uninteresting number, with interestingness determined by membership in the Online Encyclopedia of Integer Sequences (OEIS) [41]. Johnston used this to create a sequence of the smallest uninteresting numbers, a great candidate sequence for membership in OEIS.

the Platonic form of mathematical realism to those who believe that pure mathematics is nothing more than a social construct that bears no relationship to reality.[7]

Adopting a broad cognitive perspective, however, one question is why we should focus this debate on mathematics. There are lots of other cognitive phenomena whose existence is not easy to decide with our fuzzy conceptualization of existence. Examples of this include: mothers, phonemes, objects, algorithms, centers of mass, emotions, colors, beliefs, symbols, reminders, jokes, accordions, control conditions, outrage, charity, etc. In each case—including common and abstract nouns—these terms certainly are used to describe components of reality. Yet it is not clear what sense *if any* these terms might have an objective existence that is independent of our cognitive systems. There are certainly arrangements of physical matter that everyone would call an accordion, but that doesn't make accordions (or objects, or planets) a *thing* in a metaphysical sense. It's not even clear that the physical world should be described as having independent objects or parts in any objective sense.[8]

If our representation of "exists" can't decide existence for these ordinary nouns, what is it good for? It could be that the meaning of "exists" is set to be useful for ordinary human communication about present physical objects, not for deciding metaphysical questions. The following list shows typical instances of "exists" used in the spoken language section of the American National Corpus [22]:

- ... But uh, they lived in a little country town outside of Augusta called Evans, and Evans still **exists**, and that's where a lot of my relatives still, still are too ...
- ... and all the auxiliary personnel the school bureaucracy that **exists** mainly to perpetuate itself...
- ...the technology **exists** to to check it if they want...
- ... do you know if that **exists** today in a cross country kind of collective...
- ... I personally feel like as long as the Soviet Union **exists** it's going to be a threat to the United States and...
- ... died down on that I thought it was interesting that recently here the Warsaw Pact no longer **exists** as a military force but it's merely an economic now ...
- ... yeah it's the best one that **exists** I guess so ...
- ... uh the opportunity **exists** for someone to get killed ...
- ... um it's kind of amazing the disparity that **exists** between ...
- ... but I'm I'm about uh thirty miles west of there so you have uh actually green trees and such that you don't notice that that other part of New Jersey **exists** actually ...

As these examples illustrate, "exists" is typically used to talk about objects not in terms of metaphysical status but rather in concrete terms relevant to some shared assumptions by all speakers and listeners. In many cases, existence is not even asserted but *assumed* (presupposed) ("...it's kind of amazing the disparity that **exists** between..."). It is no wonder that whatever our conception of "exists" is, it is ill-suited to metaphysical discussions. The concept was acquired through usage like these—or more likely, even more simplified versions that occur in speech to children—without any sense of existence as is debated by philosophers.

Clearly what is needed to debate the existence of, say, natural numbers is a refined notion of (metaphysical) existence, one that is not confused by our ordinary commonsense notions. However, it may also be the case that our commonsense "exists" cannot be refined enough—perhaps the metaphysical "exists" is too far from our commonsense notion for us to analyze such questions.

Sameness

Other puzzles in mathematical ontology concern questions of sameness. For instance, is $f(x) = x + 4$ and $g(x) = \frac{x+4}{1}$ the same, or different? When can—or should—two mathematical objects be considered to be the same?

[7] Even more extreme, some have argued that pure mathematics is inherently sexist in its form and assumptions [23]. For discussion, see [42].

[8] This "insight" is behind a lot of new-age mumbo jumbo.

The problem here is another version of the problem encountered with existence: such questions ignore the fact that sameness is a fuzzy cognitive concept. In fact, one classic puzzle in philosophy deals with our notions of sameness for non-mathematical concepts—the problem of Theseus' ship. As the story goes, the ship has its planks removed one at a time and replaced. At what point does it stop being the same ship? A variant is the story of George Washington's axe which had its handle replaced twice and its head replaced once. The ship puzzle is actually one of practical import for the USS Constitution, which is the oldest naval vessel still working. It was commissioned by George Washington in the 1790s, but most of its parts have been replaced. At what point will it stop being the same ship?

Part of what makes sameness interesting semantically is that our conception of it is flexible. We can discuss whether or not all action movies have "the same plot" or "all my friends have the same haircut" while ignoring the obvious trivial senses in which those statements are false. This debate even occurs somewhat in cognitive science with the debate about whether learners could ever get something fundamentally "new"—is it the case that everything learners create must be a combination of what they've had before (nothing is new, everything is the same) [12], or can learners create something fundamentally unlike their early representations [4].[9]

One common response to Theseus' ship is to consider the semantics of sameness or identity ("is"): indeed, other languages distinguish notions of sameness. German, for instance, has separate words for exact identity ("das selbe") and similar but different copies ("das gleiche"). Programming languages have similar constructs (e.g., Python's "is" versus "equals"). Scheme includes at least three equivalence operators (*eq*?, *eqv*?, and *equal*?) for testing different senses of equality, in addition to numerical equality (=). For the mathematical examples f and g above, we might consider variants of sameness including sameness of intension versus extension. Given that the semantics of English are ambiguous and we could in principle distinguish the relevant meanings, why think that there's any objective truth to the matter about sameness for Theseus or mathematical objects? In all cases, the puzzle arises from using a word whose imprecision is easy to miss.

The nature of mathematics

Some of the most basic questions about mathematics concern what mathematics is and what mathematical objects are. At the very least, mathematical terms are good descriptions of neural representations in mathematicians' heads. When a mathematician says that they discovered a proof by noticing a 1-1 correspondence between integers and rational numbers, that statement surely describes some cognitive process that relates mental representations. Work in cognitive science has sought to discover the cognitive origins of many ordinary mathematical ideas, like numbers [4] and geometry [8, 43]. Work in neuroscience has identified neural signatures of approximate number (e.g., estimation) [27–30], its hallmarks in other species [2, 3, 10], and derived some core characteristics of human numerical estimation as an optimal solution to a well-formalized problem of representation [32, 35, 44]. In all of these senses, mathematical concepts are good scientific descriptions of certain types of cognitive systems, a view close to psychologism in the philosophy of mathematics.

But of course, much of the debate about mathematics is whether there is *more* than what is in mathematicians heads. Psychologism was famously argued against by Edmund Husserl [21], who emphasized the unconditional, necessary (a priori), and non-empirical nature of logical truths, in contrast to the conditional, contingent, and empirical nature of psychological truths (see also [18]). Indeed, there is a sense in which mathematical ideas do seem fundamentally non-psychological, even though they are realized primarily in computational systems like brains and computers.

Perhaps the mystery of what mathematical concepts essentially are reduces yet again to imprecision—in this case, imprecision about the concept of mathematics itself. For many, mathematics means arithmetic (math majors are constantly asked to calculate tips on restaurant bills), even though some of the best mathematicians I know were terrible at adding and subtracting. It takes experience and broad training to realize the extent of mathematical thinking—that it is not even about numbers, for instance, but about concepts and their relationships. These concepts are sometimes numbers, but only in the beginning. Advanced mathematics moves on to study the properties and relationships for creatures much more abstract than integers, including functions, spaces, operators, measures, functionals, etc. One of my favorite examples of the breadth of mathematical thinking is Douglas Hofstadter's *MU* system [20],

[9] One way to make progress in the debate is to formalize precisely what might be meant by learning, novelty, and sameness, showing for instance how learners could create particular representations in learning that are not the same as what they start with, even though the capacity to do so must be "built in" from the start (for work along these lines in numerical cognition, see [33]).

a Post canonical system with four rewrite rules: (i) a U can be added to any string ending in I, (ii) any string after an M can be doubled, (iii) Any III can be changed to U, and (iv) any UU can be removed. Hofstadter's question is: can MI be changed to MU using these rules? The solution—a simple feat of mathematical thinking—is to notice the invariance that (ii) can double the number of Is, and (iii) can reduce that number by three. Therefore, the rules leave the property of whether the number of Is is divisible by 3 unchanged, so they cannot change a string MI with one I (not divisible by 3) to one with zero Is (divisible by 3), MU.

Given the range of activities that are described as mathematics, it is not so surprising that there are a range of ideas about what mathematics is at its core. Indeed, many of the positions on the nature of mathematics seem quite reasonable. There is a sense in which mathematical concepts exist independent of minds (Platonism), a sense in which mathematical truths are empirically discovered (Empiricism), a sense in which mathematics is "just" about formal systems (Formalism), a sense in which it is about our own concepts (Psychologism), a sense in which it is a useful lie or approximation (Fictionalism), and a sense in which they are socially—or cognitively—constructed in that we could have chosen other formal systems to work with. Proponents of any view would likely argue that their version of mathematics is fundamentally incompatible with others, and that might be true.

It needn't be the case that only one of these views is correct. Like the terms discussed throughout this paper, the cognitive representation of mathematics and mathematical objects might themselves be imprecise. What is considered to be a core feature of mathematics might depend largely on the context of the comparison. When compared to physics, chemistry, or biology, a defining feature of mathematical thinking might be its focus on a-priori truth. When mathematics is compared to systems of violable rules like music, the relevant feature of mathematics might be on its stricter adherence to formalism. When compared to natural sciences, the key feature of mathematics might be that it takes place in the minds of mathematicians. The psychological underpinnings that support basic mathematical notions like integers, points, and lines might have resulted historically from our particular cognitive systems; at the same time, the insights required for mathematical progress might be most similar to art.

Such a fluidity with essential features is well-known in psychology. Amos Tversky and Itamar Gati argue for a *diagnosticity hypothesis* [47] where people may choose features to compute similarity over that are most diagnostic in clustering a set of objects. In their simple example, England and Israel are judged to be more similar when presented in the context of Syria and Iran than when presented in the context of France and Iran. The intuition behind this is that people may judge similarity based on a diagnostic feature like religion (Muslim-vs-Not) when Syria and Iran are present, and a different feature (European-vs-Not) when France is in the mix. Cognitive similarity also obeys other kinds of interesting, nontrivial properties, including being asymmetrical: North Korea is more similar to China than China is to North Korea. These studies show that people's conception of the relevant features to a concept depend considerably on context, so it may be a waste of time to argue about what that feature "really is." In this light, it is not so surprising that philosophers come to different conclusions about the core properties of mathematics and mathematical objects. The mistake is in thinking that any one of them is right.

Conclusion

The perspective I have argued for straddles a line between major mathematical positions and does not commit to strong statements about mathematical ontology, realism, or language. The reason for this is that strong positions are not justified. Strong views are only sensible when a precise question has been asked. It is useful to consider and evaluate the strong positions for whether $2 + 2 = 4$ or Goldbach's conjecture is true. It is not useful to debate the strong positions for whether a baby can be a bachelor. Natural linguistic and conceptual systems are simply not up to the task: such questions are clearly questions of definitions and we pick definitions that are useful. Even worse in cases like bachelor, we may not pick them at all, instead relying on fuzzier or stochastic notions that are so far hard for anything but human-raised-humans to acquire.

Natural language terms like "planet," "bachelor," and "mathematics" can of course be interrogated as part of cognitive science. There is a superbly interesting question about how people might represent words and what words might mean for the cognitive system. But, natural language terms cannot be interrogated usefully as part of metaphysics. The reason is that the meanings of natural language are only about the forces that shaped our cognitive science. We evolved moving far from the speed of light, on a timescale of decades, and a distance scale of a few meters. Our language and cognitive systems often intuitively capture phenomena on this level of the universe, but are hopelessly inadequate outside of this domain. For instance, we have a difficult time conceptualizing the very small or very large, or relativistic effects like time dilation that are seen closer to the speed of light. And even for things inside the domain of our evolutionary heritage, our cognitive and linguistic systems did not grow to *accurately* represent the world, but rather to *usefully* do so. Survival utility explains why we have percepts of colors in a multidimensional space (rather than a linear wavelength, like the physical space), why we like BBQ and sex, and why we feel pain and joy even though there is no pain and joy outside of our own brains. Our natural concepts were not made for truth, and cannot be trusted for truth. This insight lies behind the success of the scientific revolution.

In large part, the illusion may explain why the sciences have made remarkable progress on understanding the world but philosophical problems linger for centuries. Could it be the case that determining whether natural numbers "exist" is a harder problem that that of determining the origin of species or the descent of man? It seems, instead, that such problems are simply ill-posed or incoherent at their core, and this prevents useful progress.[10] The problem may in part be that natural language does not have the tools to tell us that its questions are impossible to answer. This contrasts with mathematics, where a nice example is provided by proofs showing that Cantor's *Continuum Hypothesis* is independent of the axioms of Zermelo-Fraenkel set theory [6, 16]. It is remarkable that mathematicians were able to resolve this deep question by showing that the assumptions did not determine an answer. Similarly strong answers cannot even be aspired to for philosophical questions that are stated without axioms, like those under consideration in mathematical ontology and areas of philosophy more generally. It is not easy to see what method might tell us that our questions are unanswerable, other than stepping back and realizing that the questions are (critically) asked in natural language, and natural language has misled us. Progress may only be made by realizing that the hypothesis under consideration is too incoherent to even be evaluated—or, in the words of Wolfgang Pauli, "not even wrong."

It may be temping to try to resolve questions of mathematical ontology or philosophy more generally by refining natural language terms. This has been the program of several approaches to philosophy of mathematics and philosophy of science. However, it is hard to see how dedication to terminological precision could end ultimately end up with anything other than a formal mathematical system. If we provided notions of existence or sameness that were grounded in a logical system, it feels as though we would be no closer to discovering the truth about the fundamental questions we thought we asked. Analogously, resolving the question of whether prime numbers exist in the context of an axiomatization of number theory is not going to tell us whether prime numbers exist according to our intuitive conceptualization of "real" (metaphysical) existence. If the psychological notions of existence are inherently imprecise, it is likely that there is no clearly definable metaphysical question lurking behind our psychological intuitions, just as there is no such question for "or" and "bachelor." Only psychological questions can be asked about these systems, not questions about the nature of reality. In this case, the appearance of a profound and usefully debatable question would only be an illusion.

Acknowledgements

I am grateful to Celeste Kidd, Ernest Davis, Julian Jara-Ettinger, Mark Frank, and Lance Rips for providing extremely helpful comments or discussion relevant to this work.

[10] Philosophical problems do sometimes get solved, although often by the natural sciences. A recent example from cognitive science is work on Molyneux's Question of whether a person who had never seen could, upon activation of their vision, visually recognize objects they had only touched. Is there an innate amodal representation of shape/space? Studies of curable blindness have provided a negative answer [19].

References

[1] Sharon Lee Armstrong, Lila R Gleitman, and Henry Gleitman, *What some concepts might not be*, Cognition **13** (1983), no. 3, 263–308.

[2] E.M. Brannon and H.S. Terrace, *Representation of the numerosities 1–9 by rhesus macaques.*, Journal of Experimental Psychology: Animal Behavior Processes **26** (2000), no. 1, 31.

[3] J.F. Cantlon and E.M. Brannon, *Basic math in monkeys and college students*, PLOS Biology **5** (2007), no. 12, e328.

[4] Susan Carey, *The Origin of Concepts*, Oxford University Press, Oxford, 2009.

[5] Rudolf Carnap, *Pseudoproblems in philosophy* (1928).

[6] Paul J Cohen, *The independence of the continuum hypothesis*, Proceedings of the National Academy of Sciences of the United States of America **50** (1963), no. 6, 1143.

[7] K Patricia Cross, *Not can, but will college teaching be improved?*, New Directions for Higher Education **1977** (1977), no. 17, 1–15.

[8] Stanislas Dehaene, Véronique Izard, Pierre Pica, and Elizabeth Spelke, *Core knowledge of geometry in an Amazonian indigene group*, Science **311** (2006), no. 5759, 381–384.

[9] Daniel C Dennett, *Intuition pumps and other tools for thinking*, WW Norton & Company, 2013.

[10] J. Emmerton, *Birds' judgments of number and quantity*, Avian visual cognition (2001).

[11] Philip M Fernbach, Todd Rogers, Craig R Fox, and Steven A Sloman, *Political extremism is supported by an illusion of understanding*, Psychological science (2013), 0956797612464058.

[12] J.A. Fodor, *The language of thought*, Harvard University Press, Cambridge, MA, 1975.

[13] J.A. Fodor and Z.W. Pylyshyn, *Connectionism and cognitive architecture: a critical analysis, Connections and symbols*, A Cognition Special Issue, S. Pinker and J. Mehler (eds.) (1988), 3–71.

[14] Jerry Fodor, *The present status of the innateness controversy* (1981).

[15] Susan A Gelman, *Psychological essentialism in children*, Trends in cognitive sciences **8** (2004), no. 9, 404–409.

[16] K Gödel, *The consistency of the continuum hypothesis*, Princeton University Press, 1940.

[17] Noah D Goodman and Daniel Lassiter, *Probabilistic semantics and pragmatics: Uncertainty in language and thought*, Handbook of Contemporary Semantic Theory, Wiley-Blackwell (2014).

[18] Robert Hanna, *Husserl's arguments against logical psychologism (prolegomena, §§) 17–61)*, Prolegomena **17** (2008), 61.

[19] Richard Held, Yuri Ostrovsky, Beatrice de Gelder, Tapan Gandhi, Suma Ganesh, Umang Mathur, and Pawan Sinha, *The newly sighted fail to match seen with felt*, Nature neuroscience **14** (2011), no. 5, 551–553.

[20] Douglas Hofstadter, *Gödel, Escher, Bach: An eternal golden braid* (1979).

[21] Edmund Husserl, *Logical investigations vol. 1.*, Routledge& Kegan Paul, 1900.

[22] Nancy Ide and Catherine Macleod, *The American National Corpus: A standardized resource of American English*, Proceedings of corpus linguistics 2001, 2001.

[23] Luce Irigaray, *Is the subject of science sexed?*, Cultural Critique **1** (1985), 73–88.

[24] Eric Margolis and Stephen Laurence, *Concepts: core readings*, The MIT Press, 1999.

[25] Douglas L Medin and Andrew Ortony, *Psychological essentialism*, Similarity and analogical reasoning (1989), 179–195.

[26] R. Montague, *English as a formal language*, Linguaggi nella societae nella tecnica (1970), 189–224.

[27] A. Nieder and S. Dehaene, *Representation of number in the brain*, Annual review of neuroscience **32** (2009), 185–208.

[28] A. Nieder, D.J. Freedman, and E.K. Miller, *Representation of the quantity of visual items in the primate prefrontal cortex*, Science **297** (2002), no. 5587, 1708–1711.

[29] A. Nieder and K. Merten, *A labeled-line code for small and large numerosities in the monkey prefrontal cortex*, The Journal of neuroscience **27** (2007), no. 22, 5986–5993.

[30] A. Nieder and E.K. Miller, *Analog numerical representations in rhesus monkeys: Evidence for parallel processing*, Journal of Cognitive Neuroscience **16** (2004), no. 5, 889–901.

[31] Judea Pearl, *Causality: models, reasoning and inference*, Vol. 29, Cambridge Univ Press, 2000.

[32] Steven T. Piantadosi, *A rational analysis of the approximate number system* (under review).

[33] Steven T. Piantadosi, J.B. Tenenbaum, and N.D Goodman, *Bootstrapping in a language of thought: a formal model of numerical concept learning*, Cognition **123** (2012), 199–217.

[34] Steven T. Piantadosi, H. Tily, and E. Gibson, *The communicative function of ambiguity in language*, Cognition **122** (2011), 280–291.

[35] RD Portugal and BF Svaiter, *Weber-Fechner Law and the Optimality of the Logarithmic Scale*, Minds and Machines **21** (2011), no. 1, 73–81.

[36] Geoffrey Pullum, *Plausible Angloid gibberishs*, 2004.

[37] Hilary Putnam, *Reductionism and the nature of psychology*, Cognition **2** (1973), no. 1, 131–146.

[38] Eleanor Rosch, *Cognitive representations of semantic categories.*, Journal of experimental psychology: General **104** (1975), no. 3, 192.

[39] Leonid Rozenblit and Frank Keil, *The misunderstood limits of folk science: An illusion of explanatory depth*, Cognitive Science **26** (2002), no. 5, 521–562.

[40] Bertrand Russell, *The philosophy of logical atomism*, The Monist (1919), 345–380.

[41] Neil James Alexander Sloane et al., *On-line encyclopedia of integer sequences*, AT & T Labs, 1999.

[42] Alan Sokal and Jean Bricmont, *Fashionable nonsense: Postmodern intellectuals' abuse of science*, Macmillan, 1999.

[43] Elizabeth Spelke, Sang Ah Lee, and Véronique Izard, *Beyond core knowledge: Natural geometry*, Cognitive Science **34** (2010), no. 5, 863–884.

[44] J.Z. Sun and V.K. Goyal, *Scalar quantization for relative error*, Data compression conference (dcc), 2011, 2011, pp. 293–302.

[45] Ola Svenson, *Are we all less risky and more skillful than our fellow drivers?*, Acta Psychologica **47** (1981), no. 2, 143–148.

[46] Hunter S Thompson, *Kingdom of fear: loathsome secrets of a star-crossed child in the final days of the american century*, Simon and Schuster, 2003.

[47] Amos Tversky and Itamar Gati, *Studies of similarity*, Cognition and categorization **1** (1978), no. 1978, 79–98.

[48] Kazunori Yoshizawa, Rodrigo L Ferreira, Yoshitaka Kamimura, and Charles Lienhard, *Female penis, male vagina, and their correlated evolution in a cave insect*, Current Biology **24** (2014), no. 9, 1006–1010.