Imagining the Infinite

Daniel Waxman

Draft version: please don't cite

Abstract

This paper investigates the procedure of *conceiving of a model* of mathematical structures, where this involves a distinctive kind of visual mathematical thinking. I argue that (i) conceiving of structures in this way is best understood as an exercise of the imagination (and not, as many philosophers would contend, an exercise of rational intuition); (ii) once the relevant kind of imaginative capacity is clarified, it becomes apparent that we can in fact conceive of infinite mathematical structures; and (iii) by doing so, we obtain justification in the consistency or coherence of certain mathematical theories.

1. Consistency, Visualisation, and Intuition

The topic of this paper is the epistemology of consistency: the question of whether, and if so, how, we are justified in believing that our best mathematical theories – paradigm examples being Peano arithmetic (PA), Zermelo Fraenkel set theory with the axiom of choice (ZFC), and so on – are consistent. It is a striking fact that so many mathematicians and philosophers who have considered the subject are strongly convinced that they are. Solomon Feferman, for example, writes "I, for one, have absolutely no doubt that PA [...] [is] consistent".[1] The leading set theorist Hugh Woodin is reported to have once offered to give up his chair to anyone who could demonstrate an inconsistency in the theory ZFC + "there exists infinitely many Woodin cardinals" – an extremely powerful theory extending the usual axioms of set theory with large cardinal axioms to the effect that certain very large sets exist.

Where does this confidence come from, and how is it justified? A very natural hypothesis is that our belief in consistency is justified, for many of the most common axiomatic theories used in mathematics, because we are able *to conceive of a model of those axioms*, where this involves a certain kind of visual mathematical thinking.[2] The main

[1] Feferman (2000), p. 72.

[2] I do not claim that this is the only possibility; in other work, I discuss a number of other potential

aim of the paper will be to understand what exactly this ability consists in, but first let me provide what I take to be a paradigmatic example of conceiving of a model. Let us focus on the case of the natural numbers, and in particular, a certain mathematical theory of them: Peano arithmetic, a theory in whose consistency we – most of us – are inclined to invest a very high level of confidence. The axioms of the theory are as follows:

PA-1 0 is a natural number;

PA-2 0 is not the successor of any natural number;

PA-3 For every natural number n, its successor $S(n)$ is a natural number;

PA-4 For all natural numbers n and m, if $S(n) = S(m)$, then $n = m$ (sucessors are unique);

Induction If $P(0)$ holds, and if, when $P(n)$ holds, then $P(n+1)$ holds, then $P(n)$ holds for all natural numbers n.[3]

What does it mean to conceive of a model of PA in the sense that will occupy us? I want to consider the following procedure, which I expect will be familiar to many. I take it that all of my readers are able – and probably have been since an early age – to visualize a *number line*. This might be done in a number of ways, differing in a number of details, but for the sake of definiteness I'll assume that it involves visualizing something like the picture below: an image with the form of a horizontal line with a leftmost mark at the origin (labelled with the numeral "0"), proceeding indefinitely rightwards with marks occurring at evenly spaced intervals (labeled with numerals denoting successive natural numbers).

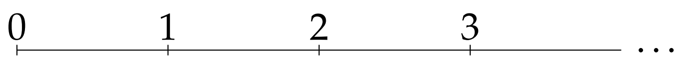

In reflecting on the process of visualizing a number line, it is hard not to think we can thereby come to obtain justification in the consistency of the Peano axioms. Why might this be? The obvious answer is that each of the axioms of PA is true when

routes. One takes justification in consistency to be inductively or abductively based (perhaps because no inconsistency has yet been discovered, or because the theories in question can be applied so fruitfully within natural science). A problem with this view is that it seems to overgenerate, by justifying the consistency of theories such as Quine's New Foundations whose consistency is plausibly an open question. Another possibility is that the consistency of our theories can be proven from other principles we accept, the most prominent candidate being that our theories are *true*. However the epistemic status of such proofs is not clear; in Waxman (ms), I argue that they are epistemically useless because they exhibit a problematic kind of epistemic circularity.

[3] For more on Peano arithmetic, see for instance Smith (2013).

suitably re-interpreted about the visualized structure. That is, suppose we re-interpret the axioms so that they concern features of marks in a number line, rather than numbers themselves, as follows:

NL-1 The leftmost mark is a mark;

NL-2 The leftmost mark is not to the right of any mark;

NL-3 For every mark, there is a mark to its right;

NL-4 For any marks, if their successors are idential then they are identical;

NL-Induction If P holds of the initial mark, and if, when P holds of a mark then P holds of its successor, then P holds for all marks.

Under this reinterpretation of the PA axioms, we can seemingly "read off" the truth of the axioms in our visualised structure. Check matters for yourself against the image you hold in mind: the leftmost mark is, indeed, surely a mark; it is not the successor of anything (for there is nothing to its left); every mark has a mark to its right (because the marks extend indefinitely rightwards); and successors are unique (because there is one and only one mark immediately to the right of any mark).

For many people, myself included, our confidence in the consistency of the axioms of PA arises from carrying out something like this procedure. We visualize a number line, consider the axioms of the theory, and notice that, suitably reinterpreted, the axioms hold true in the visualized structure. This procedure – let us call it "the visualization procedure" for PA – is a paradigm case of the epistemic procedure that will be the focus of our attention in this paper. Many philosophers have taken it, or something similar, to be of great significance in the epistemology of mathematics. Solomon Feferman, for example, traces the source of his conviction in the consistency of arithmetic to the fact that "we have an absolutely clear intuitive model in the natural numbers".[4]

My interest in this paper is not, however, primarily in the psychology or phenomenology of conviction. Even if I am right that something like the visualization procedure causes many of us to believe that arithmetic is consistent, the main question I want to focus on is epistemic: is it possible to make a case that conceiving of a model of mathematical theory in this way *rationalizes* or *justifies* our conviction that it is consistent?

The main tasks in providing any kind of satisfying answer are these: to characterize what exactly the procedure amounts to in the first place, and to discern its epistemic significance. That is, we're in the market for an account of at least the following: (i)

[4] Feferman (2000), p. 72.

the psychological processes or states involved in visualizing or possessing a conception of a mathematical structure or model; (ii) the content, broadly speaking, of these processes or states (to include any "representational" or "presentational" content, truth or satisfaction conditions, and the like); and (iii) the epistemological significance of undergoing or being in a process or state with the relevant content – in particular, whether (and if so, how) justification in *consistency* can eventuate.

I'm going to defend a particular set of answers to these questions: (i) undergoing the visualization procedure is essentially to exercise the faculty of *imagination*; (ii) we can imagine scenarios in which our best mathematical theories are *true* (or at least, true when suitably reinterpreted as being about the imagined scenario); and that imagining scenarios in this way gives us justification to believe that the relevant theories are consistent. So a lot of work needs to be done. The structure of the paper is as follows. Section 2 sketches a view of imagination and its contents; Section 3 applies that view to mathematics, and in particular addresses the problem – which faces us immediately when talking of our best mathematical theories – of imagining infinite scenarios. Section 4 turns to the epistemology of imagination, and explores ways in which the link between imagination and justification in consistency might be established. Section 5 considers and answers some objections before Section 6 concludes.

Before beginning the main work of the paper in earnest, I want to discuss a prominent alternative view that will serve as a useful stalking horse as we proceed. There is a long-standing tradition in the philosophy of mathematics that places something like *rational intuition* at the heart of its epistemology. A famous exemplar of this tradition is Gödel, who notoriously wrote (regarding set theory) that

> despite their remoteness from sense experience, we do have something like a perception also of the objects of set theory, as is seen from the fact that the axioms force themselves upon us as true. I don't see any reason why we should have less confidence in this kind of perception, i.e., in mathematical intuition, than in sense perception.[5]

In recent years, such views have undergone a resurgence in the hands of a number of philosophers.[6] Although the details vary, authors within this tradition have tended to press very hard on the analogy between intuition and perception, resulting in accounts on which intuition is viewed as a kind of faculty of "intellectual perception", presenting us with "intellectual seemings" which possess a (perception-like) "presentational phenomenology" and which are capable of providing us with insight into, and in par-

[5] Gödel (1964), p. 220. And along similar lines, discussing intuition of spatial structures: "In its purely mathematical aspect our Euclidean space intuition is perfectly correct, namely it represents correctly a certain structure existing in the realm of mathematical objects."

[6] See for instance BonJour (1998), Chudnoff (2013), and Bengson (2015).

ticular, justified or knowledgeable beliefs about, abstract states of affairs. It might naturally be thought that the visualisation procedure described above is most satisfactorily understood as an exercise of rational intuition of this kind. Such an account would presumably flesh out answers to the questions raised above as follows: (i) "visualising" a number line (as well as other examples that we would describe as "conceiving of a model of a mathematical structure or set of axioms") is fundamentally best characterized as an exercise of rational intuition, and in particular one that gives us quasi-perceptual access to a realm of abstract objects; (ii) the content of the intuition is (or includes) the *truth* of the relevant axioms (i.e. that they are true of the intuited abstract objects); and (iii) an intuition to the effect that p gives rise to *prima facie* justification that p.

However, as has been often noted, there is something deeply mysterious and unsatisfying about postulating a faculty of rational intuition that is capable of giving us access to what Tait (pejoratively) calls Models-in-the-Sky.[7] Of course, this is more an expression of dissatisfaction than an argument, and there are a few different ways that such an argument might go. Most saliently, intuitional views seem particularly badly placed when it comes to confronting the well-known Benacerraf-Field challenge, that is, the problem of explaining the reliability of the faculty of intuition *vis a vis* the target domain. The objection can be pressed a bit further with the help of a distinction due to Joshua Schechter between two different explanatory demands.[8] One is *etiological*: to explain how human beings, with our particular evolutionary etiology, could come to possess a faculty of intuition with the property of reliably generating beliefs about a realm of abstract objects. The other is *operational*: to explain the mechanisms by which such a faculty operates, in particular how reliability could be attained.The problem for the intuitional view is simply that no plausible answer seems to be available to either of these pressing questions. Naturally, this is only the beginning of an argument that would have to receive much more development to be made watertight. But that is not my aim here; I don't seek to refute the intuitional view in this paper. But I do think that these considerations make it reasonable to want to examine alternatives. And when we do, as I'll argue, we'll see that (even if the difficulties can be surmounted), we don't *need* to postulate such a faculty of intuition, at least not in order to give a compelling overall account of the psychology and epistemology of visualization procedures in mathematics. On the contrary: there is a simpler, more parsimonious, less mysterious, and phenomenologically more faithful account: our ability to visualize mathematical structures involves the exercise of the imagination. At first blush, the fact that visualization procedures can play a role in the epistemology of mathematics seems to play into the

[7]Tait (1986).

[8]Schechter (2010).

hands of theorists of rational intuition. But if the main account of this paper is along the right lines, no recourse to intuition is needed: visualization in the relevant sense can be explained by assimilating it to a cognitive faculty – the imagination – which is in no way mysterious and which we have independent reason to postulate. It is to that explanation that we now turn.

2. Imagination and its Contents

My aim in this section is to argue that the visualization procedure is an exercise of the imagination. In this section I will clarify the notion of imagination I have in mind and say something about its contents.

My focus is going to be almost exclusively on *sensory* or *perceptual* imagination. That is to say, I will not be talking about what is sometimes called "creative" – an ability to combine ideas in creative ways – or "recreative" imagination – the ability to simulate the mental states of or "put one's self in the shoes" of others.[9] Nor will I be concerned with a use of the word "imagine" that functions like "suppose" or "suppose falsely", where this involves no distinctive phenomenology.

Imagine a cat leaping at a bird, or Bruce Springsteen playing guitar, or driving on a rainy road at night. If you have complied, I take it that you will have undergone a phenomenologically distinctive imaginative episode that involved entertaining, for lack of a better term, "mental images" of various sorts – pictures in the mind's eye, or sounds in the mind's ear, so to speak. *This* is the kind of imagination that I take it is involved in our motivating cases. Imagining in this way is an occurrent mental state, whose phenomenology is very plausibly closely related to the phenomenology that would be involved in *perceiving* a cat leaping at a bird, or Bruce Springsteen playing guitar, etc. The link between perception and imagination is not coincidental. As many authors have suggested, it is plausible and helpful to think of the faculty of imagination as, roughly speaking, using "offline" the same cognitive capacities that are used online in perception.[10]

Some clarifications are in order. First, although I have been talking of "the vizualization procedure" and of "images", I don't mean especially to single out vision and the visual sensory modality. Imaginative episodes can equally well involve e.g. auditory or tactile components in addition to (or even instead of) visual ones; indeed, we

[9]See, for instance, Beaney (2005) and Currie and Ravenscroft (2002). I do not mean to deny a thesis that is commonly held by recreative theorists, namely that imagination is closely related to perception; this will be discussed further below.

[10]For a psychological account of this kind, see Kosslyn (1994). Aspects of this view of imagination are prominent in recent philosophical treatments of the imagination, for instance, Williamson (2007), Chapter 5, Yablo (2002), and Jackson (2018).

might even be able to imagine an auditory analogue of the number line (perhaps a linearly ordered sequence of bells tolling, or something of the kind). While I'll continue to speak in terms of vision and visual imagination because that is the source of the clearest examples within mathematics, I'll nevertheless propose to use the term "image" more generally than its ordinary language usage might connote, understood as covering other sensory modalities as well.

Second, an immediate objection here is worth anticipating. Aren't "mental images" dubious entities to be postulating, for are there not long-running debates within cognitive science and the philosophy of mind about whether such things even exist? Two points in response. First, I am not sure that I need to postulate mental images as self-subsisting entities; my gloss above was precisely that – a gloss – and although it's helpful to talk of images in order to pick out the distinctive phenomenology of imaginative episodes, I have no fundamental objection to doing away with such talk on ontological grounds, as long as an acceptable paraphrase is forthcoming in whatever preferred ideology replaces it. Second, I am less than convinced that there is anything dubious about postulating mental images. It's certainly true that there is a debate within cognitive science between pictoralists – those who think that in imagining, we enter into representational mental states that represent in a way closely analogous to the way that pictures do – and their opponents, who think that they represent in a non-pictoral – typically propositional or descriptional – way. But, arguably, this debate is best construed as one about the nature of (the representational features of) mental images, not about their very existence.[11]

A distinction is sometimes drawn between "objectual" and "propositional" imagination; certainly, the verb "to imagine" sometimes takes a propositional expression as a complement and sometimes a noun phrase, so that in appropriate circumstances I can be truly said to imagine either that there is a cat leaping at a bird in front of me or to imagine the cat itself. As Steven Yablo notes, the two kinds of imagining should be considered as distinct (in that one possesses "alethic" content that represents states of affairs and the other "referential" content that represents particular objects).[12] In many cases, including the ones I'll be primarily concerned with, I take it that *both* kinds of imagining are involved: we (objectually) imagine a scenario in which there's a number line, and we (propositionally) imagine that it has various properties, such as e.g. having a leftmost element. Nevertheless, in this case, the objectual imagination seems to take a certain priority. We propositionally imagine that there's a number line with a leftmost element *by way of* objectually imagining a scenario containing a number line

[11] For arguments against taking visual images to represent pictorily, see e.g. Pylyshyn (2002); for a (qualified) defense, see Block (1983).

[12] See Yablo (1993). This paper also contains a defense of a link between conceivability and (metaphysical) possibility, along the lines of those discussed in Section 4.

with a leftmost element. The proposition that there's a number line with a leftmost element is satisfied or verified by the imagined scenario; in this case, and many others, we propositionally imagine that p by imagining a scenario in which p holds.

A general thesis connecting propositional and objectual imagination might be suggested by these reflections: that whenever we imagine a scenario in which p holds, we thereby imagine that p. But consider the following situation. I successfully imagine a square in front of me; a square is an object whose group of symmetries is the dihedral group D_4; in at least one sense, it's nevertheless not the case that I imagine that there's an object in front of me whose group of symmetries is D_4 (perhaps I've never heard of the dihedral group; or perhaps I have, but I don't deploy the concept when I'm considering the imagined square).[13] In order to imagine that p by way of imagining a scenario in which p holds – at least in the sense that might plausibly do any epistemic work – we need, additionally, to *take* it that p holds in the imagined scenario. What is this extra "taking" component? Although the issue deserves further treatment, I propose to understand it as a *belief* to the effect that p holds in the imagined scenario.[14]

Let's turn now to the content of an imaginative episode at a given time. I've been talking about imagining a "scenario", and we ought to get clearer on what this involves. Here I will propose an account of the contents of imagination, drawing on work by Christopher Peacocke and Peter Kung, according to which it has several components.[15]

First, there is what we might call the *qualitative* component of an imagined scenario – the content of the pictures in your mind's eye, or the sounds in your mind's ear. Given the above-mentioned link between perception and imagination, in particular the plausibility of the conception of imagination as a kind of off-line analogue of perception towards which I am inclined, it is reasonable to suppose that such an account will exploit many of the same resources as an account of perceptual content. For the sake of concreteness, let us work with something like that in Peacocke (1992), Chapter 3. Although I'll be using Peacocke's framework, I don't think that anything especially turns on the controversial parts of his view: for instance, for Peacocke, the contents

[13] I intend here what might be called a *de dicto* sense of "imagine that...". There's plausibly a different, *de re*, sense, on which it *is* true that whenever you imagine a square, you imagine that there's an object whose group of symmetries is D_4. From the epistemological standpoint, though, this *de re* sense is less interesting, since it is *prima facie* implausible that imagining that p in this *de re* sense is itself capable of providing justification that p is consistent.

[14] I am not at all sure that the belief-based account of "taking" is ultimately correct; nevertheless, engaging more deeply with the issue would take us too far afield. Any problems here will be paralleled in the case of perception, for there are plausibly similar distinctions to be made between propositional and objectual perceiving, and between "*de re*" and "*de dicto*" propositional perception. However the "taking" component is understood, it is important to note that it must sustain a distinction between *justifiably* and *unjustifiably* taking p to hold in a given scenario. This is one reason why understanding it as belief is a natural simplifying assumption.

[15] The debt in particular is to the account of perception found in Peacocke (1992), Chapter 3, and to the account of imagination found in Kung (2010).

of perception are non-propositional, but I don't see any reason why a propositional account couldn't equally well be adapted to my purposes (supposing that it's adequate as an account of perceptual content in the first place). At any rate, on Peacocke's view, when we undergo a perceptual episode, we are presented with a spatial type, which is to say a representation of egocentric space, or way "of filling out the space around the perceiver".

A spatial type includes a number of features. There is an origin and set of axes, in relation to which the points in egocentric space are located. Furthermore, the points are presented as possessing qualitative properties: in the visual case, for instance, each point in egocentric space (distinguished as such by its distance and orientation with respect to the origin and axes) possesses features such as colour, hue, texture, saturation, and so on. So, when I perceive or imagine a grey rugby ball in front of me, the content of that perceptual or imaginative episode includes presenting me with a region of egocentric space, located to my front (in the sense that it is "frontwards", relative to the distinguished origin and set of axes), that is roughly ovoid shaped and is filled in with greyish points of various shades and hue. Imagination is, of course, distinct from perception; when we *imagine* a scenario possessing qualitative content, it is not presented to us as actual in the way that *perceiving* the same qualitative content would be. Clearly there are phenomenological differences between imagination and perception, and I don't mean to suggest otherwise the contrary. All I want to suggest here is that thinking of imagination as quasi-perceptual makes it very natural to think of imaginative episodes as having a kind of content that is structurally analogous to perceptual content, and that thinking of imagination in this way will prove to be illuminating.[16]

A second component of the content of an imaginative scenario – following a helpful discussion of Peter Kung – involves what might be called *assigned* content.[17] Imagine, if you are acquainted with it, King's College, Cambridge, on fire. Now imagine a perfect replica of King's College on fire. Even though the phenomenology and qualitative content of these two imaginative episodes are the same, the total content of the episodes surely differs between the two cases.[18] So we are led to postulate a kind of content that is the source of this difference: roughly, one that results from having different mental labels or tags assigned to or associated with the objects in the imagined scenario. In

[16]One difference that is worth flagging concerns *determinacy*: in many cases, the content of an imaginative episode will be less determinate than the content of a perceptual episode, though it is not at all obvious how this difference should best be fleshed out in a theoretical account of content.

[17]Kung (2010).

[18]The reference is to an example in Wittgenstein (1991), p. 39:"Someone says, he imagines King's College on fire. We ask him: 'How do you know that it's King's College you imagine on fire? Couldn't it be a different building, very much like it? In fact, is your imagination so absolutely exact that there might not be a dozen buildings whose representation your image could be?' – And still you say: 'There's no doubt I imagine King's College and no other building"'

the one scenario, the burning buildings are labelled <King's College>; in the other, <a replica of King's College>. Following Kung, let us call this *label* content.[19]

There is reason to think that label content is not the only kind of assigned content. Consider the following further case. You imagine [King's College on fire on a Monday]; now you imagine [King's College on fire on a Tuesday].[20] Furthermore, suppose that you imagine the very same sequence of events (the same buildings burning, the same people evacuating, the same timbers falling, etc) in both episodes. In such cases, the two imagined scenarios will have the same phenomenology, qualitative content, and label content (for both scenarios contain the same objects, labelled in the same ways). But we again want to say that they are different imagined scenarios, and thus we are led to postulate differences in their total content. This we might call "stipulated" content – part of the content of imaginative episodes that is assigned, but not due to assigning mental labels to particular objects. In the next section, we will see that the content of imaginative scenarios needs to be supplemented further. But in order to motivate that supplementation, let us first consider some complications arising for the claim that we imagine mathematical structures.

3. Imagination and the Infinite

The main thesis of this paper is that we can imagine various structures or models of mathematical theories and thereby obtain justification in their consistency. An immediate clarification needs to be made concerning the status of the "structures" or "models" here – are these supposed to be abstract objects in their own right? If so, it looks like our attempt here will be enmeshed in many of the difficult tasks that seem to confront theorists of intuition – not an appealing prospect. Fortunately, we can afford to take a fully eliminativist stance here about any such things. A helpful distinction, due to Shapiro, can be drawn between *structures*, abstract objects analogous to universals, in that they are capable of being exemplified or instantiated, and *systems,* pluralities of objects (related to one another in various ways) that are capable of exemplifying structures or satisfying axioms.[21] The thesis I defend here concerns systems, not structures. More precisely, I want to say that we are able to imagine scenarios in which there are systems of objects of various kinds – in particular, systems of objects which satisfy our best mathematical theories. There is no requirement that we are somehow able to imag-

[19] This distinction can be used to characterize the much-debated issue of whether "high level" properties are directly presented in perceptual or imaginative experience. If – as Siegel (2010) and others have argued – they are, then they will be part of *both* the qualitative and label content of the experience.

[20] The brackets here are simply to indicate scope: it is supposed to be Monday / Tuesday *in the imagined scenario*.

[21] Shapiro (1997).

ine mysterious or spooky abstract objects like structures; nothing is needed beyond an ability to imagine plain old concrete objects, as long as we can imagine enough of them and we can imagine their being related in the right ways.

But that brings us to perhaps the chief source of difficulty. The imagination-based account is supposed to give us justification in the consistency of our best mathematical theories; at least, that is the aspiration with which we began. I think it will be entirely uncontroversial that we are able to imagine some small finite systems; and indeed, perhaps that is no trivial achievement. But the theories that we've been taking as paradigms – Peano arithmetic, Zermelo-Frankel set theory, and so on – are all infinite theories in the sense that they possess no finite models. How, then, could we ever imagine systems of objects exemplifying the axioms of these theories? Isn't it plainly impossible for finite beings like us to imagine infinitely many objects in the way that would be required?

To drive the point home, consider again what happens when we try to visually imagine a number line. By doing so, we seem to imagine a model of the natural numbers, but necessarily any model of the natural numbers must contain infinitely many objects. But surely at no point during an imaginative episode can we have a visual image of *an entire model* of the natural numbers; at best, at any given time we are bound to visualize only some finite initial segment of one. If that is right – and it is surely extremely plausible – then it might seem as if the imagination-based account is simply a non-starter. In my view, this – the problem of infinity – is the most pressing problem that the account needs to face. In a moment, I'll propose my own answer – one that requires taking another look at the contents of imagination and extending it. But first, let me briefly mention a response that I'll argue is ultimately unsuccessful, proposed by Charles Parsons in a related context.[22]

3.1 Imagining Vaguely or Generally?

Parsons's chief subject is mathematical intuition. He investigates a notion of objectual intuition ("intuition of" – as opposed to propositional intuition, or "intuition that") by focusing in particular on the example of types and tokens of strings of strokes (such as e.g. |||). The relevance of this example is of course that strings of strokes, like marks in the number lines that we have been considering earlier, are structurally identical to the system of natural numbers. Now, Parsons claims that intuition can give us knowledge of the Peano axioms (suitably reinterpreted about stroke-types), and thus faces the question of how we are able to come to know claims that take the form of generalizations over *all* – infinitely many – natural numbers (/stroke types). Let us, following

[22] Parsons (1980). See Parsons (2008), Chapter 5, for further development of the view.

Parsons, focus on possible justification for the axiom of infinity here – a paradigm case of a claim which would appear to require intuition of infinitely many objects. Fully aware of this worry, Parsons offers two suggestions. The first is that the object of our intuition is *vague*, so that the relevant case is one in which we are:

> imagining vaguely, that is imagining a string of strokes without imagining its internal structure clearly enough so that one is imagining a string of n strokes for some particular n[23]

The second suggestion is, by contrast, one in which we take:

> as paradigm a string (which now might be perceived rather than imagined) of a particular number of strokes, in which case one must be able to see the irrelevance of this internal structure, so that in fact it plays the same role as the vague imagining.[24]

I take it that the ideas are closely related. Transposing the discussion back over to the idiom of imagination, both ideas can be viewed attempts to overcome the problem of infinity by taking a *particular* imagined string-type to be typical, in a way that sustains universal generalization. If such an attempt works, then claims that can be verified to hold about the particular string in the imagined scenario can be justifiably extrapolated to claims about *all* strings. The two approaches he mentions are, in effect, two different possible bases from which the generalization might be drawn. On the first approach, the content of the imaginative episode is a string of strokes with a *vague* or *indeterminate* length: in other words, a string of the form $|||| \cdots |$, where the "..." indicates that the internal structure of the string is not articulated in thought, except (presumably) insofar as it is recognizable as a string of strokes of the requisite kind. The question "how many strokes are in the imagined string?" is thus misguided – there is simply no determinate answer, because the internal structure of the string is not fleshed out enough in the imaginative episode for there to be a particular number n such that the string consists of n strokes. Nevertheless, despite the string's indeterminate length, the suggestion is that we are able to imagine extending it by adding a further stroke.

There are a few objections that might be raised for this proposal: plausibly there are difficulties in the very notion of imagining a string of indeterminate length, or in extending it in imagination (after all, it is not obvious that we are entitled to think that (i) the initial string and (ii) the extended string have distinct lengths – at least, not without some account of how something like penumbral connections are preserved in vague imagination).[25] But the more fundamental worry, as I see it, is one that is

[23]Parsons (1980), p. 156.
[24]Parsons (1980), p. 157.
[25]See Page (1993) for a development of worries along these lines.

shared with the second proposal; consequently I will treat them together. According to that second proposal, the imagined string (indeed, as Parsons notes, a perceived string would do equally well) is in fact fully determinate in thought; but its precise internal structure (i.e. the exact number of strokes it contains) is somehow seen to be *irrelevant* to the fact that it can be extended by adding another string.

Here then is the pressing objection for both proposals: how can it be that we're justified in moving from a particular claim about the indeterminate or representative string – for instance, that it can be extended by adding another stroke – to a *general* claim – that *every* string can be extended in this way? The situation is reminiscent of attempted justifications of diagrammatic reasoning in mathematics. Suppose we exhibit a particular diagram of a right angled triangle, and persuade ourselves (say) that Pythagoras's Theorem holds of that triangle; the danger of extrapolating the claim to *all* right angled triangles is of course that, as Marcus Giaquinto puts it, diagrams may "tempt one to make unwarranted generalizations, as one's thinking may too easily depend in an unnoticed way on a feature represented in the diagram that is not common to all members of the class one is thinking about"; and indeed, this worry has led many to discount the role of diagrammatic reasoning in mathematics altogether.[26] It is true that sometimes this worry can be allayed. For instance, if we reason about a particular triangle depicted in a diagram and then come to appreciate that no step in that reasoning relied on any attributes that are not shared with all members of the relevant class, then perhaps the generalization is justified. But it is hard to see how that structure is present in the case of extrapolation from imagined strings. The conviction that any string can be extended by adding a further stroke seems to be a claim that we 'read off' from the imaginative episode – it is not plausibly the outcome of any procedure of *reasoning*, let alone one that takes as a premise the 'typicality' of particular string directly imagined. Consequently it is very much doubtful whether we are justified in generalizing, on that basis, to *all* strings.

3.2 The Dynamics of Imagination

I turn now to my preferred way of handling the problem of infinity. To do so, we will have to return to and supplement the account of intuition given in Section 2. According to the story so far, imaginative content consists of qualitative content and assigned content (including labelling content). The need for supplementation arises because, for all that we've said up till now, the content of an imagined scenario might well be totally static. In particular, something needs to be added to handle the *dynamic* character of imagination. What I mean by this is best seen by considering a couple of examples of

[26] Giaquinto (2007), p. 77.

ways in which imagination and imaginative episodes can be dynamic.

(I) We can imagine scenarios within which time progresses (and, possibly but not necessarily, changes occur over time). Consider again the example of imagining a cat leaping into the air. There is a way of imagining such a scenario whose content is surely not a static qualitative visual image (even perhaps supplemented with background stipulations to the effect that the cat is jumping); rather it is something more like a *sequence* of such images, where the qualitative content of each static visual image is, in a roughly continuous manner, related to the content of the previous and next images in the sequence, so that the overall impression is of viewing an *event* of a cat leaping into the air as opposed to a mere static image of a time-slice of some such event.

(II) We can imagine scenarios from different spatial perspectives. Suppose again that you imagine a cat leaping into the air; now imagine observing the same event, but from a spatial perspective located on the other side of the cat. Notice that we are able to imagine scenarios from different perspectives in this way even if the scenario is regarded as "static", i.e. we are not imagining time as progressing within it. For instance, imagine a game of cricket is taking place, but frozen at the exact moment at which the ball is being struck by the batsman. We can imagine such a scenario from the spatial perspective of the bowler; or from the perspective of the batsman; or an audience member in the stands; or indeed from a dynamically and continuously changing perspective, for instance, as if from a camera moving smoothly around the perimeter of the stadium.

An example of Peacocke's proves instructive to work through in this connection.[27] He asks us to imagine a scenario containing a cat fully obscured by a suitcase. I think that there are actually at least three possible ways in which we might comply with that request. One – I take it that this is the way in which Peacocke himself is thinking of the case – can be stated in purely static terms, by including the presence of a cat in the scenario via a piece of stipulative content. The relevant imaginative episode then has content along the following lines: qualitative visual content including a suitcase-shaped region of visual space, coloured and shaded in an appropriately suitcase-like manner, in addition to assigned content which includes a stipulation that behind the suitcase lies a cat. The "presence" of the cat in the content of this imaginative episode – what makes it the case that we are imagining a cat obscured by a suitcase, as opposed to just a suitcase – is wholly handled by the stipulation.

By contrast, I think that there is a second, distinct way of imagining such a scenario, involving imagining it from different perspectives. Start in roughly the same way, with qualitative visual content of a suitcase-shaped region in the field of view. But then go on to imagine the scenario from a different perspective – say, from a perspective

[27] Peacocke (1985).

located from the side of the suitcase as opposed to the initial side, in which the cat *is* indeed visible. Thus the content of this part of the imaginative episode will include an explicit visual representation of a cat (perhaps, say, a black-and-white, sitting-cat-shaped region).

Indeed, there's a third way too: start with the same qualitative content again; but this time, imagine that the scenario develops as time progresses, with (say) a person entering into view and removing the suitcase, revealing the cat behind it. Presumably this way would require additional stipulated content to the effect that the cat did not just spring into existence when the suitcase was removed, that it was there all along. But it is nevertheless a way of imagining a scenario in which a cat is fully obscured by a suitcase.

The point of all of this is that it shows a convincing need to be able to theoretically distinguish dynamic imaginative episodes, of the kinds above, from sequences of static imaginative episodes. In particular, I think this shows that imaginative episodes can have richer content than that of any single static imaginative episode or even sequence of such static episodes. After all, there is a difference between (i) undergoing a dynamic imaginative episode whose content is a cricket match, first from one perspective, and then from another; and (ii) two distinct static imaginative episodes of different cricket matches, which share the same qualitative and phenomenological content as in (i). In slogan form: *sequentially imagining is not the same thing as imagining in sequence*. Let us describe imaginative episodes as *extended* when they include a dynamic element in this way.

What I want to suggest is that this aspect of imagination – the fact that we are able to have extended imaginative episodes – is best understood *dispositionally*. When we have an extended imaginative episode as of a cat jumping into the air, as in (I) above, this has two components. The first is our disposition to generate a sequence of static images that are continuously related to one another, such that each of them is as of a "snapshot" of a cat jumping into the air; and the second component is our disposition to regard these "snapshots" as being part of the *same* imaginative scenario, as opposed to different static images corresponding to a number of distinct and disconnected scenarios. The story runs along similar lines when we have an extended imaginative episode that represents us as taking, say, two different perspectives on the same cricket match. We have a disposition to generate a sequence of static images that are continuously related to one another, such that each of them is as of a "snapshot" of a cricket match; and we have a disposition to regard these "snapshots" as being part of the *same* imaginative scenario. The upshot is that the content of extended imaginative episodes are informed by dispositions two-fold: on the one hand, dispositions concerning the way in which the scenario is (going to be) dynamically elaborated, and on the other, dispositions to "glue" the various parts of the elaboration by regarding them together into a single imaginative episode.

3.3 Imagining a Model of Arithmetic

Once we understand the difference between static and extended imaginative episodes, I think we will see that the problem of infinity is not so insuperable after all. In short: once we recognize that the content of imagination can go beyond what is imagined in a single moment, it becomes at least in principle possible to see how we might be able to imagine infinite scenarios. And indeed, I think that this is not just a mere possibility – it is one that is plausibly realized when we attempt to imagine a number line.

To see this, consider again what happens when we imagine a number line. We begin with an image of a finite initial segment: a static imaginative image whose qualitative and assigned content is something like the image:

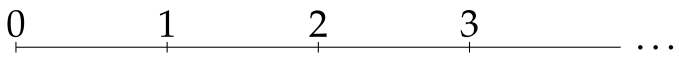

Perhaps the numerals "1", "2", etc, are part of the qualitative content of the episode; or perhaps they are assigned as labels. But either way, the crucial point is that – if your conception of a number line is anything like mine – this static image does not exhaust the content of the imaginative episode. For in addition to the static image, we have a number of dispositions to develop that image dynamically in various ways – dispositions to dynamically develop our perspective on the scenario in various ways. As we consider what would happen as we "pan our mental camera right", so to speak, we have a disposition to add to the qualitative and assigned content of the imaginative scenario: in particular, to extend the number line so that our image encompasses more strokes/numbers than before. While it is true that any given manifestation of these dispositions to "extend" the number line will still result in a finite image, the crucial point is that we have *additional* dispositions to extend that finite image even further, going on in the same way, indefinitely.[28] In this way, it is reasonable to describe the imagined number line as infinite: not because at any point we have a static image of containing infinitely many objects, but because no matter how far rightward our "mental camera" pans, so to speak, we are disposed to continue updating our perspective on

[28] As the talk of "going on in the same way" suggests, the proposal will at some point need to confront the 'rule-following considerations' of Wittgenstein (1953) and Kripke (1982). A full discussion would take us too far afield, but in short, I am optimistic that a dispositional account of some kind can be made to work. In particular, I think the apparatus in the recent defense of dispositionalism due to Warren (2018) is very promising in this regard: the composite dispositions discussed there seem not only able to resolve Kripkensteinian worries, but also to be precisely what is needed to explain the kinds of infinitary dispositions involved in the present account.

the imagined scenario by adding new strokes/numerals indefinitely; at no point will the rightward scanning ever lead to an image where the number line can be seen to *terminate*. So, even if these dispositions are never manifested – and, given the finitude of human cognition, it is hard to see how they could be – they are nevertheless relevant to the content of the imaginative episode in question; and in particular, they make it the case that the content of that imaginative episode is a scenario containing infinitely many objects.

The development of an account of imagination capable of allowing infinite scenarios has, so far, mostly taken place by reflection on our imaginative capacities. But, with that said, I think it is worth emphasizing the interesting and potentially fruitful convergence between the account presented above and some recent work by Marcus Giaquinto on visual thinking in mathematics in which he presents an account of cognition of mathematical structure.[29] In his discussion of infinite structures, he argues that we can cognize an infinite structure – for instance, a number line of the kind I discuss – by means of distinguishing between category specifications, on the one hand, which are: "feature descriptions stored more or less permanently", and, on the other, visual images, "fleeting pattern[s] of activity in a specialized visual buffer, produced by activation of a stored category specification".[30]

Giaquinto's notion of a category specification is inspired by the related notion of "category pattern" formulated by Kosslyn. As Kosslyn (1994, 74), puts it, "A visual mental image is a pattern of activation in the visual buffer that is not caused by immediate sensory input. The components of the protomodel have the same properties when they are used to "inspect" an object that is visualized instead of perceived. For example, if asked whether Snoopy the dog has pointed ears, people report visualizing the cartoon and "looking" at the dog's ears. The shape of the ears is matched to stored shapes in the ventral system, and the size, location, and orientation of the ears is registered by the dorsal system— just as would occur if the dog were being perceived."[31]

He then goes on to argue that we have a category specification that allows us to cognize infinite sets of the kind just mentioned:

> "My suggestion is that our grasp of this structured set, the well-ordered set of evenly spaced marks on an endless horizontal line, issues (or can issue) from a stored visual category specification. But this is not achieved by "reading off" the descriptions of the category specification, since we have no direct access to those descriptions. Rather, as a result of having the cate-

[29]Giaquinto (2007), Chapter 11.

[30]Giaquinto (2007), p. 108.

[31]Again, we see here the link between visual (and more generally, sensory) perception and sensory imagination.

gory specification, we have a number of dispositions which, taken together, indicate the nature of the structured set it represents. These are dispositions to answer to certain questions one way rather than another. For example: Given any two marks, must one precede the other? Yes. Do the intermark spaces vary in length? No. Is the precedence of marks transitive? Yes. Can any (non-initial) mark be reached from the initial mark by scanning to the right at a constant speed? Yes. But some questions will have no answer: Is the intermark length more than a centimetre? These answers tell us something about the nature of the mental number line as determined by the features specified in the category specification.[32]

There are certainly differences between Giaquinto's account and mine. While I have been discussing sensory imagination and the possibility of imagining infinite scenarios, Giaquinto's focus is primarily on *de re* knowledge of abstract objects such as structured sets; indeed, his language of 'grasping' and (objectually) 'knowing' abstract structures in which his discussion is formulated suggests that he may well view his account as closer to the intuitional views described in Section 1 than to the account above. Nevertheless, there is I think a fundamental similarity between his core insight and my discussion of the content of imagination: namely, that once it is understood that our faculties of imagination/cognition must extend beyond mere static images (in my terminology; "visual images" in his) and encompass extended imaginative episodes (in my terminology; "category specifications" in his) that are specified in part by our dispositions – then the problem of infinite structures no longer seems impossible to meet. In short, we have good reason to think that we can imagine the infinite.

4. Epistemological Payoff and Objections

4.1 Imagination as a Guide to Consistency

If what I've said so far is correct, we are able to imagine systems of objects satisfying at least some infinitary mathematical theories. It's time to discuss the epistemological upshot of this fact. As previously advertised, I would like to claim that if we're able to imagine a system exemplifying some mathematical theory, then we have at least *prima facie* justification to believe in the consistency of that theory. That is, I think there is good reason to accept the following principle:

Imaginability-Consistency Link if S imagines that ϕ, then S has *prima facie*, defeasible justification to believe $\Diamond_L \phi$ (where "$\Diamond_L$" expresses consistency or logical possibil-

[32] Giaquinto (2007), p. 227.

ity).[33]

Two brief comments about **ICL**, before we turn to discussing its motivation. Firstly, the notion of justification employed is intended to be *propositional*, not *doxastic*. Imagining a scenario in which ϕ holds *opens up a potential route* to forming a justified belief that $\Diamond_L\phi$; but it doesn't (for all that **ICL** tells us, and rightly) entail that *any* old belief that $\Diamond_L\phi$ held in the presence of an imaginative episode of the right kind is justified (– if, for instance, such a belief were formed on the basis of wishful thinking). Secondly, the notion of justification employed is intended to be *prima facie* and defeasible. **ICL** does not rule out – and rightly so – that any of the justification obtained from an imaginative episode of the right kind can be undermined or otherwise defeated, so that it is at least theoretically possible that someone who imagines a scenario in which ϕ holds can nevertheless fail to have any residual justification in $\Diamond_L\phi$.

What, then, can be said in favour of **ICL**? The first thing to note is that it is a member of a more general species of evidential links between conceivability and possibility: principles according to which various species of 'conceivability'-like mental states give rise to justification or evidence for possibility-claims. Such links have been much discussed: the epistemic role of the imagination, and in particular its ability to deliver justification or knowledge about non-actual states of affairs, has been heavily emphasized in recent years.[34]

My strategy for defending **ICL** will not be to attempt a full blown account of the epistemology of imagination and modality; that would be a large undertaking, and would take us far afield from the issues with which we began. Instead, I'll briefly examine three of the most promising approaches to modal epistemology – approaches which have primarily aimed at vindicating links between conceivability and *metaphysical* possibility – and will that each, if successful, very naturally carries over to **ICL**. In fact I think a stronger conclusion is warranted: because the modality figuring in **ICL** is consistency – i.e. it is broadly *logical* in nature – as we'll see, if anything our principle is more plausible than, and able to evade some of the most serious objections to, more traditional connections that have been claimed between conceivability and metaphysical possibility.

The first way that **ICL** might be defended is by appeal to *modal rationalism* of the sort defended by David Chalmers in a recently influential discussion. Chalmers distinguishes a number of dimensions of conceivability, including between (i) *prima facie*

[33]Understanding consistency as a species of logical possibility, and thus construing it as a modal sentential operator is quite natural: for a prominent defense, see Field (1984). Following Field, I will slide between talk of consistency and logical possibility. That being said, nothing I say here is incompatible with understanding consistency as a predicate of sentences instead, if that is preferred.

[34]See for instance Williamson (2007), and the essays collected in Kind and Kung (2016) and Macpherson and Dorsch (2018) for a small sampling.

(“conceivable for [a] subject on first appearances”) and *ideal* (“conceivable on ideal rational reflection”) conceivability[35]; between (ii) *positive* (“one can form some sort of positive conception of a situation in which S is the case”) and *negative* (“S is not ruled out a priori, or when there is no (apparent) contradiction in S”) conceivability[36]; and between (iii) *primary* (“it is conceivable that S is actually the case”) and *secondary* (“S conceivably might have been the case”) conceivability.[37] The notion of imaginability as discussed earlier is I think clearly characterized, in Chalmers’ terms, as a species of *prima facie*, positive, and primary conceivability.

Now, Chalmers advances the following thesis as the “most plausible entailment between conceivability and [metaphysical] possibility”:

Modal Rationalism Ideal primary positive conceivability entails [metaphysical] possibility.

The details of Chalmers’ argument are intricate, turning on subtle issues in the philosophy of content, and will not be explored in detail here. But supposing that Modal Rationalism is correct, **ICL** very plausibly follows. Notice first that although Modal Rationalism is in the first instance a claim about *entailment* relations between conceivability and possibility rather than *evidential* relations, it nevertheless clearly motivates an evidential claim. If ideal primary positive conceivability *entails* primary possibility, and if imagination is a species of *prima facie* positive conceivability, then it is very plausible to think that imagining a scenario in which ϕ gives one *prima facie* justification to think that ϕ is metaphysically possibly true. A helpful analogy might be drawn with the validity of proof, which also admits a distinction between ideal (“valid on ideal rational reflection”) and *prima facie* (“valid on first appearances”). An ideally valid proof that ϕ entails that ϕ; but possession of a *prima facie* valid proof that ϕ gives *prima facie* justification that ϕ. In both cases, the resulting justification will presumably be susceptible to undermining or other kinds of defeat by anything that defeats the *ideal* validity / conceivability of the proposition in question; but in both cases, it is hard to deny that there will be resulting justification.

So Modal Rationalism motivates an epistemic link between conceivability and metaphysical possibility. But on any plausible account of metaphysical possibility, it is a stronger notion than consistency or logical possibility – logical falsehoods are metaphysically impossible. So anyone who accepts Modal Rationalism ought to accept **ICL** also.

In passing, it is worth pointing out that that **ICL** is if anything *more* secure than its analogue concerning metaphysical possibility, because it is invulnerable to one major

[35]Chalmers (2002, p.147).
[36]Chalmers (2002, p.149).
[37]Chalmers (2002, p.156).

source of objection to Modal Rationalism.[38] Many proponents of a post-*Naming and Necessity* Kripkean conception of metaphysical modality regard it as intimately connected with the natures or essences of particular things.[39] Philosophers in this camp are often suspicious of moving from conceivability to possibility-judgements, since they believe that we may often seem to conceive of situations in which the essences of things are misrepresented (e.g. a scenario in which water is not H20 but is XYZ). However, it is very hard to see how such worries could arise for consistency or *logical* possibility: even if a situation in which water is XYZ is metaphysically impossible, it is hard to see how it is inherently contradictory. Logical possibility, unlike metaphysical possibility, is unconcerned with the essences of things.

The second approach to modal epistemology that I believe motivates **ICL** involves an appeal to a kind of *phenomenal dogmatism* in the sense recently discussed by authors such as James Pryor and Michael Huemer.[40] Although phenomenal dogmatist views have been developed most extensively in the philosophy of perception, such views can be naturally extended to motivate **ICL**, conditional on a couple of plausible (albeit contestable) assumptions about the nature of imagination, to be mentioned below.

Dogmatists about perception proceed on the assumption that perceptual states have representational content.[41] They also accept the following thesis linking this representational content with justification:

Dogmatism if one is in a perceptual state which represents that p, then one has *prima facie* justification that p.

Phenomenal dogmatists about perception hold, additionally, that this justification arises in virtue of the phenomenology of perceptual states. As James Pryor puts it, such states have "a distinctive phenomenology: the feeling of seeming to ascertain that a given proposition is true" – and, according to phenomenal dogmatists, it is this phenomenological feature of perception that underwrites the resulting justification.[42]

If it is correct, then this approach would appear to carry over relatively straightforwardly to imagination. Something would need to be said about the phenomenology of imagination, by analogy with the phenomenology of perception, in order to draw the desired conclusion about justification. But there is a very natural hypothesis here: that just as perception presents us with a seeming that a given proposition is *true*, that imagination presents us with a seeming that a given proposition is *possibly* true. And

[38] See for instance Yablo (2009).

[39] Kripke (1980).

[40] See for instance Pryor (2000) and Huemer (2007).

[41] This doesn't obviously entail that perceptual states must directly possess *propositional* content; something like the proposal could go through as long as whatever non-propositional content they have is nevertheless representational.

[42] Pryor (2004, p.357).

as long as the sense of possibility that figures in this claim is one that entails logical possibility, then it seems that there is a clear dogmatist case to be made for **ICL**.

The final approach to modal epistemology is empiricist in flavour, and emphasizes the conception of imagination, mentioned early, as using "offline" the same cognitive capacities that are used online in perception. The basic idea is that the imagination *simulates* having a certain perceptual experience; just as a perceptual episode would (in ordinary circumstances) provide us with justification that various propositions are actually true, the thought goes, an offline simulation of a perceptual episode will thus provide us with justification that various propositions are *possibly* true.[43] As before, if this approach is successful, then it is very easy to see how it carries over to justification in *consistency*, for consistency is again implied by metaphysical possibility.

No doubt, much more could and should be said about these approaches. But let me end the section by emphasizing that none of these lines of argument are intended to convince someone who antecedently rejects any epistemological connection between imagination and logical possibility – someone who we might, reasonably enough, view as somewhat akin to a sceptic.[44] James Pryor distinguishes between two different anti-sceptical projects: the *ambitious* project is to refute the sceptic on his or her own terms, appealing to resources he or she would accept.[45] By contrast, the *modest* anti-sceptical project attempts to establish to *our* satisfaction – we who are not antecedently sceptical – that scepticism is unwarranted in this domain and that the anti-sceptical view forms a coherent and stable platform that is not susceptible to internal critique, even if this platform would not be accepted by a die-hard sceptic. I take the options discussed in this section to be an attempt at a sketch of how the modest project might go. At the very least, we ought to have optimism that a link from imaginability to logical possibility can be defended as part of a coherent, stable, and plausible package of views.

5. Objections and Replies

Objection: Your view justifies too much. Take, for example, naive set theory – a classic example of a theory endorsed by various mathematicians and philosophers (some tacitly, some – like Frege – explicitly), but which turned out to be inconsistent. Didn't these people believe themselves to have a conception of a universe of sets satisfying the axioms of naive set theory? And if so, aren't you committed to saying they were justified in believing the theory to be consistent?

[43] Views of this kind are defended in Peacocke (1985) and Jackson (2018).

[44] This character would be interestingly analogous to someone who denied that perceptual experiences were epistemically connected to the formation of justified belief or knowledge about the *actual* world.

[45] Pryor (2000).

Response: If this description of the set theory case is correct, I think my view delivers the correct verdict. To begin with, I don't think there's anything especially damaging about accepting the conclusion of the objection here. I've been explicit that any justification in the consistency of axioms that arises from imagination is defeasible; one incontrovertible source of defeat would be a proof of a contradiction from the axioms in question, such as Russell provided for naive set theory. I think there would be nothing unusual or in any way problematic about saying that a number of mathematicians and philosophers *were* justified in believing that this theory was consistent, only for this justification later to be undermined. We have learned to live without Cartesian certainty everywhere else; why think that it is somehow required when dealing with issues of consistency?

The qualification at the beginning of the last paragraph is important, however – for it is not at all clear that the naive conception of set was ever taken to be justified on this basis, i.e. our ability to imagine a structure in the sense that has been discussed above. It is commonly noted that there is a distinction between two different concepts of *set* or *class*: (i) a logical notion – roughly, the extension of a predicate – and (ii) the familiar mathematical notion – roughly, objects that lie somewhere in the cumulative hierarchy obtained by starting with the empty set and proceeding into the transfinite via a powerset operation. I suspect that those who believed in (the consistency of) naive set theory tended to have the first concept of set in mind, and that they were not imagining a scenario in which they took the axioms to be true.[46] If that is right, then, obviously enough, no justification for the consistency of the theory could be forthcoming along the lines discussed in this paper.

Objection: No matter how things stand with naive set theory, there still plenty of other examples of cases where your view delivers the wrong verdict. Consider Escher drawings which seem to represent spatially impossible scenarios. What, on your view, is stopping us from imagining such scenarios, and thereby (wrongly) gaining justification in their consistency?

Response: First, the point about the defeasibility of justification recurs, and I think in itself takes much of the sting out of the objection.

Second, while Escher-style drawings (and analogous imaginative episodes) may well represent spatially impossible scenarios, it is not at all obvious that these scenarios are *logically* impossible. Here, as before, the difference between metaphysical possibility and logical possibility may prove to be significant in undermining certain potential counterexamples to the link between conceivability and logical possibility.

Third, in such cases, there is a good case to be made that we are *not* in fact successfully imagining scenarios with the (spatially inconsistent) properties that they appear

[46] See e.g. Parsons (1976) for evidence to this effect in the case of Frege.

Waterfall (M. C. Escher, 1961)

to have. This is a point which can be seen by again appealing to the distinction between static images and extended imaginative episodes. It is true that we are able to imagine Escher-style scenarios in the sense of possessing a *static image* that seems to have, e.g., an always-descending waterfall, in it. But if (as will be clear if the reader attempts to do so) one attempts to dynamically elaborate this static image by considering the imagined scenario from a different perspective on the imagined scenario, it soon becomes obvious that this cannot consistently be done – there is no way of, for instance, consistently imagining what the *back* of the waterfall in the picture would look like. So it is not clear that we can have extended imaginative episodes of such paradoxical scenarios, and in fact (on reflection, at least) clear that we cannot. Of course, we may justifiably but mistakenly *think* that we can have such extended episodes; and possibly even get justification in consistency on that basis; but such justification would be very easily defeated.

Objection: Doesn't your view *overgenerate* justification in a way that potentially crowds out the epistemic role of proof in mathematics? Surely when we imagine scenarios modeling various mathematical theories, we don't just take them to verify the *axioms* of mathematical theories, but also other claims as well. For instance, grant that we imagine a model of the natural numbers, and thereby propositionally imagine the truth of the axioms of PA (suitably reinterpreted). What is stopping us saying that we also propositionally imagine, in addition, elementary claims about the natural numbers – say, that every number is even or odd – that are usually thought to require proof?

Response: There are a few things to say in reply. First, one source of difference that is certainly epistemically relevant in the context of a communal endeavour like mathematics is that proofs are communicable, whereas imagination is private. So at the very least, there is *something* to be gained by giving a proof, in that the result is now shareable in a way that a mere description of an imaginative episode is not.

Second, there's a difference in the content of the justification: the epistemic pay-off of proving a result is not the same as successfully reading it off in an imagined scenario. Reading off a claim gets you justification in its consistency (perhaps in conjunction with other claims also read off in the same scenario) – continuing the example from the objection, if $\mathsf{PA} \wedge E$ is the theory adding to PA the claim that every number is even or odd, the justification obtained by reading off the claims in this way will be in something like the proposition that $\Diamond_L(\mathsf{PA} \wedge E)$. But a *proof* of the fact that every number is even or odd will establish the very different claim $\Box_L(\mathsf{PA} \rightarrow E)$, or equivalently $\neg\Diamond_L(\mathsf{PA} \wedge \neg E)$. This means that the role of proof is in no way crowded out; imagination at best gets us justification that some claim is mutually consistent with the basic axioms governing a structure, but a proof of a claim from axioms gets us justification that it *necessarily* follows, in the sense just mentioned.

Third, it's true that on my view we can obtain justification in $\Diamond_L(\mathsf{PA} \wedge E)$ in two different ways: one that proceeds by way of reading off the axioms and E from an imagined scenario; and one that proceeds inferentially, by inferring $\Diamond_L(\mathsf{PA} \wedge E)$ from (i) $\Box_L(\mathsf{PA} \rightarrow E)$ – whose justification comes from giving a proof of E from the PA axioms – and (ii) $\Diamond_L\mathsf{PA}$ – which, let us suppose, is antecedently justified via imagination. But although the same claim can be justified via both of these methods, that is not to say that there is no epistemic value to be had from using the proof-based method. For in many ways and for many cases, giving a proof from axioms will be a much safer way of obtaining justification in the claim $\Diamond_L(\mathsf{PA} \wedge E)$ or its analogues. One way to see this is just to consider the potential defeaters for either of these methods. How might the justification generated by running through a proof be defeated? Well, by (among other things) making a mistake in the proof, or making a mistake in the inference from the premises to the conclusion, or by anything that defeats initial justification in the axioms. But by contrast, the imagination-based method will be defeated by anything that jeopardizes the reading off of E from the imagined scenario. Now perhaps in the case of very simple claims, like the example of E that we have been considering, the imaginative procedure does not seem so risky. But this is a result of the simplicity of E – when, on the other hand, we are dealing with more complex claims, or claims that stretch our abilities of visualization, it ought to be clear why, on the present view, the possession of a proof is of epistemic worth.

6. Concluding Remarks

If all that I've said above is correct, then the imagination-based account I've sketched above provides a very nice treatment of what it is, and why it is epistemically important, to conceive of a model of a mathematical theory. Furthermore, it gives us an account of a central type of justification – what might even be described as a form of logical or mathematical intuition – in a way that doesn't require the postulation of any mysterious faculties.

There is, however, a remaining source of dissatisfaction worth mentioning and responding to. If what I've said is right, then imagination provides a route to explaining and rationalizing our confidence in the consistency of our best mathematical theories. But this might be thought to be a relatively paltry payoff in comparison to the intuitional account, which seems to offer, by contrast, justification in the *truth* of those theories. There are two things to say in response, however. The first is that if the imagination-based account really does point the way to justification in consistency, then that is no mean feat. In light of Gödel's incompleteness theorems, no consistent theory of a certain minimal strength is able to prove its own consistency. If we are indeed justified in believing in the consistency of our best theories, then that justification must either come from proving it within some other theory, or by some means outside of mathematical proof altogether. The imagination-based account, I belief, elegantly resolves this issue.

The second point to make in response is that there is a striking trend in the philosophy of mathematics according to which consistency (and closely related notions like coherence and logical possibility) are central to the epistemology of the subject. A large number of prominent views – including structuralism, formalism, plenitudinous platonism, abstractionism, and perhaps even fictionalism – contend, roughly speaking, that the consistency or coherence of a theory is fundamentally all that is required in order for it to legitimately serve as the basis for mathematical inquiry.[47] If a view of this kind is on the right track, then the imagination-based account is capable of playing a key role in the justification of much ordinary mathematical theorizing.

I'll close by emphasizing a final virtue of the account and indicating some further work to be done. The virtue is simply that, if my account's correct, it goes some way to explaining the importance of a *conception* of mathematical subject-matters. For example, much has been written about the concept of set, and in particular, the iterative conception of sets as the members of a cumulative hierarchy beginning with the empty set and being generated in stages via the powerset operation.[48] It is, I think, a common attitude that this work is relevant, not just to the anthropological issue of unpacking *our* concept of set, but to the epistemological issue of *justifying our acceptance* of set-theoretic axioms.[49] But the relevance of exploring our conception of set for this latter project is not at all clear. If what I have said is along the right lines, then one explanation suggests itself: our possession of a conception of some mathematical subject-matter is closely linked to our being in a position to justifiably believe in the consistency of the axioms governing it. The question for future work is whether, and if so, to what extent, this approach can be applied to justification in the consistency of *set theory itself* – whether, for instance, we can reasonably be said to be able to imagine the cumulative hierarchy of sets; and if not, how large of a fragment of the axioms of ZFC *can* be justified on this basis. But that is a large topic, and work for another day.

References

Balaguer, M. (1998). *Platonism and Anti-platonism in Mathematics*. Oxford University Press.

[47]See for instance Shapiro (1997), Hellman (1989), Weir (2010), Balaguer (1998), Hale and Wright (2001), and Field (1980).

[48]See, for instance, Boolos (1971).

[49]This is explicit in, e.g., Paseau (2007).

Beaney, M. (2005). *Imagination and Creativity*. Open University.

Bengson, J. (2015). The Intellectual Given. *Mind 124*(495), 707–760.

Block, N. (1983). The Photographic Fallacy in the Debate about Mental Imagery. *Noûs 17*(November), 651–62.

BonJour, L. (1998). *In Defense of Pure Reason*. Cambridge University Press.

Boolos, G. (1971). The Iterative Conception of Set. *Journal of Philosophy 68*(8), 215–231.

Chalmers, D. J. (2002). Does Conceivability Entail Possibility? In T. S. Gendler and J. Hawthorne (Eds.), *Conceivability and Possibility*, pp. 145–200. Oxford University Press.

Chudnoff, E. (2013). *Intuition*. Oxford University Press.

Currie, G. and I. Ravenscroft (2002). *Recreative Minds: Imagination in Philosophy and Psychology*. Oxford University Press.

Feferman, S. (2000). Does Reductive Proof Theory Have a Viable Rationale? *Erkenntnis 53*(1-2), 63–96.

Field, H. (1980). *Science Without Numbers*. Princeton University Press.

Field, H. (1984). Is Mathematical Knowledge Just Logical Knowledge? *Philosophical Review 93*(4), 509–552.

Giaquinto, M. (2007). *Visual Thinking in Mathematics: An Epistemological Study*. Oxford University Press.

Gödel, K. (1964). What is Cantor's Continuum Problem. In H. P. P. Benacerraf (Ed.), *Philosophy of Mathematics*, pp. 116–117. Prentice-Hall.

Hale, B. and C. Wright (2001). *The Reason's Proper Study: Essays Towards a Neo-Fregean Philosophy of Mathematics*. Oxford University Press.

Hellman, G. (1989). *Mathematics without Numbers: Towards a Modal-structural Interpretation*. Oxford University Press.

Huemer, M. (2007). Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research 74*(1), 30–55.

Jackson, M. B. (2018). Justification by Imagination. In F. Macpherson and F. Dorsch (Eds.), *Perceptual Imagination and Perceptual Memory*, pp. 209–226. Oxford University Press.

Kind, A. and P. Kung (2016). *Knowledge through Imagination*. Oxford University Press UK.

Kosslyn, S. M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. MIT Press.

Kripke, S. (1980). *Naming and Necessity*. Harvard University Press.

Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language*. Harvard University Press.

Kung, P. (2010). Imagining as a Guide to Possibility. *Philosophy and Phenomenological Research 81*(3), 620–663.

Macpherson, F. and F. Dorsch (2018). *Perceptual Imagination and Perceptual Memory*. Oxford: Oxford University Press.

Page, J. (1993). Parsons on Mathematical Intuition. *Mind 102*(406), 223–232.

Parsons, C. (1976). Some Remarks on Frege's Conception of Extension. In M. Schirn (Ed.), *Studien zu Frege*, pp. 265–277. Friedrich Frommann, Stuttgar.

Parsons, C. (1980). Mathematical Intuition. In *Proceedings of the Aristotelian Society*, Volume 80, pp. 145–168. JSTOR.

Parsons, C. (2008). *Mathematical Thought and Its Objects*. Cambridge University Press.

Paseau, A. (2007). Boolos on the Justification of Set Theory. *Philosophia Mathematica 15*(1), 30–53.

Peacocke, C. (1985). Imagination, Experience, and Possibility. In J. Foster and H. Robinson (Eds.), *Essays on Berkeley: A Tercentennial Celebration*. Oxford University Press.

Peacocke, C. (1992). *A Study of Concepts*. MIT Press.

Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs 34*(4), 517–549.

Pryor, J. (2004). What's Wrong with Moore's Argument? *Philosophical Issues 14*(1), 349–378.

Pylyshyn, Z. W. (2002). Mental Imagery: In Search of a Theory. *Behavioral and Brain Sciences 25*(2), 157–182.

Schechter, J. (2010). The Reliability Challenge and the Epistemology of Logic. *Philosophical Perspectives 24*(1), 437–464.

Shapiro, S. (1997). *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press.

Siegel, S. (2010). *The Contents of Visual Experience*. Oxford University Press USA.

Smith, P. (2013). *An Introduction to Gödel's Theorems*. Cambridge University Press.

Tait, W. W. (1986). Truth and Proof: The Platonism of Mathematics. *Synthese 69*(3), 341–370.

Warren, J. (2018). Killing Kripkenstein's Monster.

Waxman, D. (ms). Truth, Transmission Failure, and the Justification of Consistency.

Weir, A. (2010). *Truth Through Proof: A Formalist Foundation for Mathematics*. Oxford University Press.

Williamson, T. (2007). *The Philosophy of Philosophy*. Blackwell.

Wittgenstein, L. (1953). *Philosophical Investigations*. Macmillan.

Wittgenstein, L. (1991). *The Blue and Brown Books: Preliminary Studies for the 'Philosophical Investigation'*. Wiley-Blackwell.

Yablo, S. (1993). Is Conceivability a Guide to Possibility? *Philosophy and Phenomenological Research 53*(1), 1–42.

Yablo, S. (2002). Coulda, Woulda, Shoulda. In T. S. Gendler and J. Hawthorne (Eds.), *Conceivability and Possibility*, pp. 441–492. Oxford University Press.

Yablo, S. (2009). No Fool's Cold: Notes on Illusions of Possibility. In Oup (Ed.), *Thoughts*. Oxford University Press.