learning by doing: mathematical knowledge, truth, causation

Adam Morton

fairly polished DRAFT.

> One cannot invent the structure of an object. The most we can do is to patiently bring it to the light of day, with humility — in making it known it is "discovered". If there is some sort of inventiveness in this work, and if it happens that we find ourselves the maker or indefatigable builder, we aren't in any sense "making" or "building" these structures. They hardly waited for us to find them in order to exist, exactly as they are!
>
> Alexandre Grothendieck *Récoltes et Semailles*

Physical human animals know a sliver from the seam of abstract mathematical fact. This prompts the "PB reaction", from Plato to Benacerraf and beyond: "but these domains are so different".[1] The reaction prompts a rejoinder "no, they aren't", which can take the form of assimilating mind to the abstract (Plato and Platonism down to our day) or of assimilating mathematics to the physical (20th century formalism and its contemporary continuations). Neither is necessary. We can understand mathematical knowledge and the ways it is acquired in a generally causal way appropriate to thoroughly causal beings, while understanding mathematical fact in a way that is neutral about its causal status.[2] A contemporary understanding of knowledge is consistent with a traditional understanding of mathematics as describing an abstract non-physical domain. It is also compatible with a more naturalistic picture of mathematics. This paper shows how the combinations are possible.

If knowledge were justified true belief the task would be different and more straightforward. An all-purpose account of truth could be glued to a requirement of rationality and "all" that would be needed would be an account of mathematical rationality. Proof would be at its centre, as evidence is at the centre of empirical belief. But knowledge has become a more causal concept and we now understand rationality as

1 Plato *Meno* 80-86, *Phaedo* 72-3, *Theaetetus 189*, Benacerraf (1973), Steiner (1973), Reznik (1975), Jubien(1977), Hale and Wright (2002).

2 Taking the worry at face value can prompt drastic claims, as in Callard (2018). I am not arguing that any such claims are wrong, simply that they are not needed.

a diverse and diffuse business with links to causal processes that result in true belief. So the task of linking abstract objects and particular brains faces more obstacles. My claim, though, is that all-purpose accounts of knowledge and truth are available and indeed natural, and have no incompatibility with a physical conception of human minds and their thinking. In fact, a causal connection between mathematical facts and our beliefs about them emerges as a non-mysterious natural matter.

The idea is somewhat subtle but straightforward. A rough motivation can be given with an image, so let us begin with that.

feeling truths

We often learn as a result of what we can and cannot do. This is clearest not with vision but with touch, where sensation and manipulation, as the word suggests, are inextricable. If you reach in the dark to feel the shape of an object you press it with your fingers and experience its contours so that by some combination of these you emerge with information about its solidity and outline. Or suppose you are finding out how many coins there are in your pocket. You reach in and find that you can put your first, second, and third fingers on the coins, while the thumb roams around and finds no more. Accomplishing this physical act of counting tells you that the coins form a three-membered set. (They and your fingers satsfy Hume's principle, that equinumerous sets are those that can be put into one-to-one correspondence; but you know this only by moving the fingers.[3]) The crucial difference is that the basis for the belief that is formed is information about which actions have succeeded. It is not itself evidence about the manipulated objects, although features of these objects shape tactile success and failure.

Another example, closer to what happens in mathematics, is seeing whether a plane figure is a projection of a three-dimensional shape by trying mentally to move its vertices towards or away from one. There is no new perceptual contact with the figure, just the

3 Hale and Wright (2002) discuss Hume's principle in connection Benacerraf's problem, leaving fingers out of it. There is impressive evidence that numerical thinking uses brain regions associated with the fingers — Soylu (2019), Buijsman (2019), Penner-Wilger and Anderson (2013) — and thus that counting by acts of touching and pointing is deeply connected with numerosity. Fingers appear in Kant's discussion of arithmetic (*Critique of Pure Reason* B 15–16).

awareness that an attempt has succeeded or failed. Whether or not we classify the activity as perceptual, it is carried out in causal isolation from the figure and gives information about it because mental activities have particular results.[4]

When we explore mathematical possibilities we handle them by trying out various lines of argument, various images and simplifications, various definitions and symbolizations. Seeing where these lead, including the times when they lead nowhere, tells us how the topics constrain our thinking. The constraints are the mathematical facts, a point to return to. Besides the link between basic arithmetic and the fingers, elementary mathematical thinking involves the ability to summon and shape spatial imagery, and measuring, gathering, or rotating in space.[5] When things become even a little bit sophisticated it is important to find dedicated quasi-linguistic representations of them, and symbolic activity becomes important.[6] All this needs more detail, but for now the point is the way that we can learn as a result of what we can and cannot *do*.

the unity of knowledge

There are many ways of learning something mathematica. You can prove it from first principles; you can derive it from a powerful general theorem (and both of these can be done both meticulously and with some helpful handwaving); you can take it on authority from a trusted book or teacher; it can just seem so obvious to you that you would not to try to get it from anything else; you may be convinced from pondering over a diagram. All of these, and more, are familiar parts of our histories. And all of them can go wrong even when it seems that everything is in order. Each employs some capacities that are not exploited in others. So each will have connections with thoughts that we do not usually think of as mathematical. One complicated and revealing case is when a mathematical result with a diversity of applications slowly emerges from a line of thinking in physical science. (Besides the emergence of geometry from terrestrial surveying,

4 Kanamori (2018) discusses aspect perception in connection with mathematics. His interest there, however, is not with mathematical knowledge but with the effects of conceptual point of view on the force of a concept.

5 Dehaene and others (1999), Dehaene (2009). Margolis (2019) gives evidence that the apprehension of even small integers draws on a dedicated number sense.

6 Lewis (1991) is admirable for its presentation of mathematics in plain English. But, enticing though this is, it's harder that way. See also Morton(1996).

which left echoes up to the twentieth century, and the present mathematical stimulus of string theory, examples where the transition was clear and eventually completed were the emergence of the calculus of variations from extremal principles in physics and the emergence of abstract probability theory from discussions of practical decisions and empirical evidence.) There is then a pull to coherence in two directions, to mathematically related thoughts and to procedures and beliefs of the nonmathematical kind.[7] This often connects what we take to be mathematical knowledge with what we take to be knowledge but not in mathematics. Since we call both of them knowledge and since we are often unsure which side of the line we are on, we cannot think that knowing in mathematics is radically different from knowing anything else.

The unity cannot consist in procedurally correct proof, since it is not associated with all sources of knowledge within mathematics, let alone those in the borderland. Proof must be a particularly important means to some more general end. And considerations that show that some shatterproof procedures will lead reliably to true results must get their significance from the links with the more general ideas of reliability and truth. In the last quarter of the previous century ideas about how to characterize the dependence of a belief on a reliable process began with a simple emphasis on deliberate processes which for solid causal reasons tend towards true results, as in Goldman (1976). These demand more control than is intuitively required, tend towards asking that knowledge lead to knowledge of knowledge, which also seems too strong, and are hard to fit to semi-automatic modes of everyday awareness. The next, still generally causal, accounts replaced this with conditions of causal attunement, often expressed in counterfactual terms. If the facts had been different the result would have been a different belief (as in Nozick, discussed below). These were succeeded by current anti-luck formulations, often with a flavor of safety (as in Sosa and Prichard, discussed below). The original intention was always to formulate a condition requiring the right kind of causal connection between the fact and a resulting belief.[8] The belief is true *because* the facts are as they are. It is

7 Quine (1981 and elsewhere) famously argues that the scientific pull generates the mathematical one. There is an enormous literature on this, helpfully surveyed in Collyvan (2019).

8 “Knowledge first” approaches, inspired by Williamson (2000), refrain from defing knowledge with conditions on belief. That is consistent with stating non-definitional principles linking knowledge and belief. On a more drastic approach, also suggested by Williamson, one would avoid the concept of belief altogether. My discussion mentions “beliefs”, “conclusions” and other representations, as a way of leaving the issue open, and hinting that a deeper discussion may need attention to the bearers of knowledge. I suspect that one can talk simply of states that show the appropriate sensitivity to possible variations on

possible, however, to attempt to reformulate the machinery in non-causal terms. My claim is that this is not necessary, because there are natural conditions with plausible causal interpretations which are generally congruent with the contemporary orthodoxy and apply as well in mathematics as elsewhere. In mathematics too a belief can be held because the facts are as they are.

The link between the various kinds of knowledge, if contemporary epistemology is on anything like the right track, is thus the generally causal connection between thinking something and its being true. Various forms of causation can make the connection: the direct causal relation between events, causal laws of nature, and counterfactual (subjunctive) conditionals are all possible forms. In all of them, and others, there is a general becauseness that can connect the results of a variety of inquiries to the intended objects and properties.[9]

A central reason why we care about knowledge is that it indicates how secure a person's grasp of the truth is, and, roughly correlated with this, how well suited to produce further truths the origins of a particular belief are.[10] It is hard to see how a non-causal account of knowledge could deliver this. However while truth is a binary business the truth-producing capacity of the thinking behind a belief comes in degrees and can be measured in many ways. When we say that someone, for example, knows as a result of their teacher's instruction that the derivative of sine is cosine we are discussing not simply the relation between the person and the fact of basic calculus but implicitly the route she took to her acquaintance with it and the varied capacities of that route. It is a reasonable assumption that causally nearby situations often resemble those that are likely to occur, so that more truth among possible situations at a time translates roughly into more truth among actual situations over time. I doubt that we would have much use for the concept of knowledge if this assumption were very misleading. Learning that the teacher's grasp

actual situations. Not now.

9 Knowing how can take a place here beside knowing that, since it reflects facts about which actions will have which effects. It is also related to causation in the opposite direction, from the agent to the world, which is significant for the general theme of this paper.

10 Why knowledge as an indicator of reliability? There is a strand in the philosophy of science which dismisses the concept of knowledge as a rough commonsense fudge. I think that philosophers who take this line underestimate how flexible knowledge is, how tunable to different needs, as this paper illustrates. Talk simply of causal connections between thoughts and objects is not adequate here: we want the right kinds of causal connections, and modal epistemology supplies them. Still, if there were another way my general strategy should apply to it also.

of basic calculus is shaky would suggest that their instruction on this topic is likely to be flawed and thus that although the student ended up with a true belief in this case in similar cases there may well be trouble, how frequently and how different the cases depending on how stringent our criteria for knowledge are.[11] So unless our standards are lax the student's true belief falls short of mathematical knowledge. But they should too demanding either, since they should fit real people well enough that accessing a particular person's grasp of a particular fact will tell us something useful about their future inquiries and their value as a source of testimony.[12]

Are higher standards appropriate to mathematical knowledge? The question will recur throughout this article. On the one hand knowledge represents a desideratum, something we would like as much of as possible, other things being equal. On the other hand what is possible, or feasible, or worth the effort, is usually limited, and varies from case to case and topic to topic. So applications of the concept must find a trade-off between the two. But there is no inevitable balance-point for this trade-off, and the question is the uniformity of the balance between achievable aim and adequate accomplishment.

sensitivity and safety

There are two basic strategies for describing fact/belief connections, and they differ for necessary truths such as those of mathematics. To see the contrast between them compare two formulas, both counterfactual conditionals giving simplistic criteria for a proposition to be known by a person, that loyalty to the two strategies might suggest. (Once we move beyond these simple conditionals the contrast between the underlying approaches is less stark, as we will see. But the contrast is still revealing.) The sensitivity conditional requires that if the facts had been otherwise the belief would not have been held. *If ~p then ~B(p)*. The safety conditional requires that if the person had formed an incompatible belief then the facts would have been different. *If B(~p) then ~p*.[13] They

11 Most of almost anyone's mathematical knowledge depends in part on testimony and instruction. This is true of all kinds of knowledge, but there has not been much attention to the mathematical case.

12 Welbourne (1986) and Craig (1990) describe the role of knowledge-ascriptions as evaluations of testimony.

13 The grandparents of these formulations are Robert Nozick and Ernest Sosa, though neither they nor those they influenced insisted on anything so simplistic. Nozick (1983), part 3, Sosa (1999). Note that Nozick uses a somewhat non-standard conditional, and the analysis is nearer to a biconditional than would appear from this and many expositions. See also Greco (2016), Pritchard (2018), Rabinowitz

are not equivalent essentially because the counterfactual does not contrapose, but this is a front for a deeper difference between them.[14] Both ask for a match between the reasons that the person has the belief and the facts that make its content hold, but the sensitivity conditional looks first at the facts and requires that the belief be connected to them while the safety conditional looks first at the belief and requires that the facts be connected to it. Neither insists on a causal connection running in just one way. Safety is not to be taken as a spooky influence of thought over things, but as a more flexible description of the target. The intention is better presented as "she would only have believed that it was false if it *was* false", or "if she had believed the opposite it would have been because the opposite was the case". (If she had believed that the minimum of the function was at x = 3, instead of at x= 0, it would have been because the function in question did have its minimum at 3. Compare "if the damp match had lit (it would have been because) the air was unusually oxygenated".) Similarly, sensitivity need not require one-way influence of fact on belief. The reason that one would not have had the belief if things had been different might be that things would have been different because of one's variant belief, as in some cases of self-fulfilling self-knowledge. (If the lever were in a different position she would have put it there and would have been aware of what she had done.)

Still sticking with the simple parodistic formulas, sensitivity runs into problems with necessary objects of knowledge. "If the square root of 144 were not twelve then p" is either unevaluable or trivially true for any p. Safety also has problems: "if p then the square root of 144 would not be twelve" is unevaluable or trivially false for any true p.[15] But in this case there is an appealing fix. Suppose that someone is idly calculating square roots and comes to the conclusion that the root of 144 is twelve. While still operating as a fine arithmetician she could easily have chosen a different calculation, and arrived instead at the result that the root of 169 is thirteen, or that the cube root of 1728 is twelve. She would not have concluded that that the square root of 144 is 11 or that the

(internet).

14 For the failure of contraposition see Bennett (2003, pages 143-5), building particularly on Lewis (1973). Bennett's own account of counterfactuals is particularly relevant to the present project because it minimizes the difference between regular and "backtracking" conditionals (page 208).

15 The counterfactual is applicable in one case. Suppose that if a person had failed to believe a claim it would have been because its terms would have meant something else, which would have been true. (For example, she would not have agreed to "5+2 = 7" if the numerals were interpreted base 5.) Then "if not belief then content not true" is okay I shall ignore this complication in what follows.

cube root of 1728 is 12.3. Whatever conclusion in some relevant range she had arrived at, because of her capacity for arithmetic it would have been a true one.

For dealing with knowledge of necessities a focus on alternatives to beliefs as well as on their negations is valuable. Then weaknesses in a person's grasp of a truth will show up as both contingent and necessary falsehoods that she might have arrived at. The vital question becomes the range of alternatives appropriate to a particular person in a particular context when arriving at a particular conclusion.[16] But before grappling with this consider the kind of link between fact and belief that is suggested. It is causal in an interesting way. We are dealing with a causal process that results in a belief and whose further outline can be described by specifying what other conclusions it would lead to under other conditions. Any explanation of how this occurs will depend in an essential way on the truth of the belief. (Put rhetorically: Try explaining why someone believes root 144 to be 12 on the assumption that root 144 is *not* 12. But see below.) So in this sense part of the reason why the person has the belief is the fact that makes it true, plus other closely related facts. The activity that results in the belief has a certain shape, which reflects the fact which makes the content of the belief true. It is like the way that when you are adding integers by counting on your fingers the way you touch your fingers is an image of the sequence of the integers themselves. I shall return both to the formal resemblance and to the metaphor.

Another reason for considering alternative conclusions that the person might well have arrived at is that the actual conclusion may be correct and moreover established by the person's reasons, for example by a solid proof, but still be a dubious case of knowledge. This will happen when the person's grasp of the solid proof they have produced is wobbly; they could easily have come up with a fallacious one without realizing it. (A rank amateur “solves” an outstanding problem in number theory, which seems to attract rank amateurs. The solution is good but there are many fallacious variants that he would have been as content with.) Similarly, someone could believe a mathematical truth on the basis of a proof that was in fact a valid proof of something entirely different. Moreover someone could be in this situation not only in actuality but in near alternatives to it.

16 These are different from the alternatives considered in the "relevant alternatives" approach to knowledge. For one thing, they represent variant states of the person rather than variant objective possibilities.

I draw a couple of morals from these examples. First, that we should distinguish between two kinds of conclusion. There are conclusions that a person would draw in nearby circumstances as a result of different lines of thinking, for example the beliefs that different deceptive teachers, or the same teachers on a different whim, might give. And contrasting with this there are conclusions that a person would be led to by the same line of thinking if circumstances were different, as they would be if for example the person were in a hastier or more cautious mood. Alternatives of the second kind are the more significant ones in assessing the epistemic status of a belief, at any rate a mathematical belief. Second, in both of these cases we should consider more than just a few alternative beliefs. We should consider enough possible but not too remote alternatives that we have a real grip on the person's tendencies to true conclusions.

Some of the difficulty of describing causal relations between mathematical necessities and contingent facts such as people's beliefs will also hold between milder necessities such as laws of nature and particular contingent events. Your belief in universal gravitation is surely a result of gravitation itself, among many other things, but the fact that gravitation is universal plays a different role than, say, your thinking about the evidence. It is a condition under which causes of contingent events such as the acquisition of the belief itself can take place. In order to countenance the conditions for a particular thought we do not have to consider mind-boggling counterfactuals such as "if there was not a gravitational force between all massive objects then I would not be having this thought". The move from event causes to background conditions is part of a strategy to sidestep this.

I shall consider epistemic and semantic theories where the major role is played by cause-enabling conditions, rather than causal relations between events. They too stand or fall with the way the world actually works, and they too engage with the physical mechanisms of thought and how it connects with the environment. Of course, no category of causality alone is enough to establish knowledge. A thoroughly unreliable belief-acquisition can be an effect of the fact that makes it true together with a cause which operates because of this fact, and similarly for more sophisticated causal accounts.

(Action can also operate by enabling rather than directly causing. One can be responsible for a coin's landing heads to the extent that one got it out of one's pocket and tossed it so that chance could operate. Not realizing this might make our control of random events — the most contingent of all — seem as mysterious as our knowledge of necessary facts.[17])

kinds and seriousness of exceptions

The topic is mathematical knowledge rather than knowledge in general. But some of what has been labelled as mathematics belongs to the same category of less than ultimate necessity as some of what does not get so labelled, and any sharp demarcation will be rather arbitrary. Examples are results from before the late nineteenth century separation of pure mathematics from abstract physics. Moreover, it is essential that mathematical knowledge emerge as real knowledge, continuous with other kinds, not as the result of a deliberate gerrymander. So we need at least an outline of knowledge that includes a full range of its sub-species, and allows them to merge into one another.

A person has a belief obtained by the use of a method in a particular situation. The combination could approach divine omniscience by hyper-strong perfect counterfactual correlation; the person would have a belief if and only if it is or would be true. That would give both the primordial sensitivity and safety formulas. It would be both perfectly safe and perfectly sensitive, but is not an ambition for the likes of us when the conditionals cover any serious range of worlds.[18] For one thing, the result would not serve many of the purposes of knowledge, such as evaluating the credentials of informants. There is a range of ways we can weaken the force to a human scale while balancing between sensitivity and safety. At one end of the range there are conditions appropriate to successful apprehension of the immediate environment using our powerful but limited and inflexible capacities. The apprehension could easily not have happened and a reasonable aim is that it track the surroundings accurately. A simple sensitivity

---

17 I have argued for a thoroughgoing duality between knowledge and accomplishment in Morton (2012), (2013).

18 Any candidate for knowledge is both true and believed in actuality, so that when the range of worlds is limitingly narrow the biconditional will be within the reach of human knowers, since vacuously true. But the resulting concept of “knowledge” won’t be very useful or very close to standard requirements..

conditional will do this, requiring that when there are real or possible slight differences in the facts there will be a change in belief, if only to agnosticism. But this surely asks for too much. The most aware person in the world will have blind spots; there will be scattered variations that she will not notice. It also asks for too little. It counts as knowledge a capacity that can handle only the very simplest departures from the way things are, as when a person could detect an object had it been 1 mm from where it is but not if it had been any further. So some room for exceptions must be built into this.

At the other end of the range lie mathematical and other conclusions that result from a lot of thinking. Then the belief-forming processes are much more flexible. They can lead to a much more varied output but the range of facts that they can detect reliably is limited and exception ridden. Mathematicians may be able to resolve some puzzle when n = 1, 2, 3 , 4 but find it impossible for n = 5, 6. (We can prove Goldbach's conjecture for n up to $4x10^{18}$ so this gives an example with n a multiple of 10^{18}.) Mathematicians a century or more later might understand why completely different techniques were needed for the recalcitrant cases. So it would not be a sensible requirement that we reliably produce true conclusions when the facts are even somewhat different. Instead, the primary aim should be to use our thinking capacities accurately, avoiding errors. Now a simple safety conditional is a beginning candidate, requiring that the conclusions that a person could easily have arrived at be true, even though there may be possibilities very near to them that correspond to none of the person's beliefs. Again this is unrealistically demanding. The best mathematician in the world can have a few propensities to slip up even on the way to perfectly solid and well understood proofs.[19] So again we should allow carefully placed exceptions. But we should place them so that they still require the agent's reasoning to have a pretty powerful falsity-avoiding capacity. There will always be many alternative histories where a person who actually arrived impeccably at a true conclusion would have ended up with a different and false one, but some of them are more of a disqualification to counting the actual belief as knowledge.

The aim, then, is to place exceptions in the unrealistic perfect counterfactual correlation so as to make it suitable for real human knowledge of various kinds. This will mean

19 There is a long list of slips by the immortals, including Lebesque and Cauchy. I don't know anywhere they are collected in one place. To complicate the picture, there are defective proofs butressed by powerful other considerations.

tuning the exceptions so as to accommodate *human* knowledge of different topics. And doing that will mean distinguishing between ways in which a person's possible situation might differ from the actual one: in the method used, the belief she has arrived at, and the situation she is in. There are various combinations of these factors. Things will inevitably go wrong for merely human inquiry, often for no fault of the inquirer; the task is to separate forgivable possible mishaps from seriously disqualifying ones.

It is bad when an inquiry results in a false belief, and methods that result in unacceptably many falsehoods are obviously to be avoided. But there are trivial falsehoods, such as thinking that there are 3,000 hairs on one's left arm when in fact there are 3,237. A crude measure of triviality is the range of possible worlds where the negation of a belief is true (one might call this the degree of necessity of its content). When a content holds in no possible world then it is very much not to be believed, although there may be no easy way of avoiding it. More generally, when the content of one false belief holds in a proper subset of the worlds where the content of a second belief holds, then arriving at the first belief would be a more serious error than arriving at the second. It is also bad news when someone could easily have gone wrong, while errors that are in a more remote neighbourhood of actual performance are less worrying.

So starting with perfect superhuman correlation between propensity to believe and the truth of the resulting beliefs, there are two broad kinds of exceptions that we should permit in some amounts, and each of these two kinds comes in varying degrees. The broad kinds are errors made by following the method which was actually used though in other possible situations, and those made by following a method that would be prompted by a possible situation. Errors of both kinds are more serious when the range of possibilities is greater, that is, when the mistaken content fails to hold in a greater range of situations or when the process that leads to it occurs in a greater range. These need to be weighed against one another. The generally less serious category is errors that might result from a remote variant situation, and the least serious of these are errors which are more remote from actuality or which produce a result which is false in fewer possible worlds. (Or we might say "with a greater potentiality to lead to falsehood", which is metaphysically less contentious but runs separable factors together.) And among *these* the least serious are where either the the contrary fact or the belief in it is located at modally remote situations. The most serious category is, correspondingly, when the actual method could easily have led to a false conclusion, even more so when that conclusion is necessarily false, even even more so when it is an ultimate necessity such as a logical principle rather than say something unchangeable which depends on contingent events (for example an identity between contingent individuals). Less serious, though potentially disqualifying, is when the person could easily have used a method which could easily have led to a falsehood, and among these cases the least serious is when the falsehood is a pure contingency, especially a description of a random event. Between these extremes there is a gradation of significance for disqualification. Charting this gradation precisely would mean an explicit weighting or trade-off between the numbers of potential cases in the relevant categories, degrees of remoteness from actuality, and the person's actual versus possible methods of belief acquisition.[20]

Knowledge has a characteristic specificity. The sources of a particular belief are linked to the truth of a particular proposition. When the emphasis is more on sensitivity it is the proposition bears more of the weight of specificity. The specific fit is best indicated by varying it. When the emphasis is more on safety the belief bears more: its variariation is the better index. In both cases the point is the specificity, though it is most easily and clearly described differently in different cases.

I doubt that there is a precise and objective way of making the trade-offs. They have to remain rough and qualitative. But this fuzziness allows a single description of knowledge of different domains. It also allows stronger and weaker interpretations of knowledge, consistent with effects of context, ideology, and individual limitation. (It does not easily yield a *measure* of strength of interpretations, though. The most we can say is that one is generally stronger or weaker than another.) But even taking the criteria of seriousness as a rough guide, they give mathematical knowledge a safety-biased orientation, where the range of contents that beliefs obtained by the method actually used is dominant. For the processes that lead to mathematical conclusions when bungled will often lead to drastic impossibilities; the person is usually modelling their thinking on patterns that are

20 These will usually be infinite, so comparisons of size are problematic. Note how beliefs based on thinking that typically leads to necessities are automatically directed towards safety-like conditions, and more contingency-directed beliefs towards sensitivity.

directed at necessary truths, and is usually choosing between alternatives in vocabulary that produce either necessary or impossible claims.[21]

This description of mathematical knowledge supports the comparison with a certain kind of knowing by doing. Suppose that you are wrapping a malleable plastic substance onto an immovable solid shape, in the dark. (Immovable by your efforts, that is.) You are going to estimate the contours of the solid by the ways you can wrap the squishable stuff. We want to detect flaws in your estimating as revealed by particular estimates. But whatever you do the solid is going to stay the same shape, so if your capacity is flawed this will not be fairly judged by discrepancies between possible shapes of the solid and shapes you could have given the plastic. Instead, flaws must be primarily marked by discrepancies between shapes you can give the plastic and those of the underlying contours. This is in the direction of safety, resembling mathematical knowledge.[22]

The solid shape corresponds to the necessary truths. (Which for all you know might be affected by greater powers than yours, something early modern philosophers worried about.) The resulting shape of the plastic matter corresponds to the representation of such a truth in your mind. (And for all you know the solid may have many features that the sensitivity of the plastic misses.) And your awareness of this shape is what gives you a clue about the shape that has formed it. But the comparison raises a question: how does the effect of the solid shape on the plastic shape come into the metaphor? There is a literal version of the question. How are we to describe the causal role of background conditions, especially if we want to avoid possible situations where these conditions are different. I discuss this below, but some preparation is needed first.

a problem about firmness of belief

---

21 Another consequence of this fact is that the identification and individuation of the method used is less problematic than it is, say, for beliefs obtained by perception or memory. Just as well, then, that comparison beliefs obtained with whatever means the person finds themselves using in nearby circumstances play a larger role in those cases.

22 Another example is prescientific knowledge of humanly short spans of time based on the regular counting one can do in them. Here too there is a connection with Kant, identifying "inner sense" with one's knowledge of one's own thought and action (*CPR* A33/B49).

If all knowledge is basically the same should not the same sort of evidence apply whatever the claim? Shouldn't we be content with inductive evidence for claims about numbers? Or, going in the other direction, might we deny that Euclid knew what he claimed to, since his stated axioms are mathematically inadequate? The attitude to knowledge described in the previous section suppports a response to these questions. First consider some examples.[23]

(A1) There are many conjectures, of which the best known may be Goldbach's conjecture and the Riemann hypothesis, asserting that all integers have some property, where we can prove that they all do up to but not beyond some large number. We do not have counterexamples for any n, but cannot prove or disprove the conjecture for the infinitely many remaining cases. We routinely say that we do not know whether these conjectures are true.[24] But there is ample inductive evidence for them, at least as much as we would need for many claims in the rest of science.

(A2) There are also many conjectures which hold in many cases, without known counterexamples, which are for various reasons much more plausible than their contraries, which although we are unable to prove or disprove them we are pretty sure are provable if true. Possibly the best known of these is the P≠NP conjecture in computer science.[25] A conjecture like this might evoke as much confidence, and be as coherent with established theories for which there is extensive hard-to-dispute evidence, as many uncontroversial scientific claims. So why are these not known if true?

(B) Sometimes we do not recognize a theorem for what it is, taking it as a law of nature or an accidental generalization. A 19th-century physicist could have seen abundant evidence that conserved quantities were associated with symmetries in the Lagrangians

---

23 Kitcher (1980) discusses cases like this and arrives at different conclusions based, I think, on assuming a fixed standard for knowledge rather than, as here, one that is sensitive to the particular modal properties of the proposition in question, as is standard in various ways in most accounts of knowledge since that time. I fall into a similar trap in ch. 3 of Morton (2003).

24 J.E. Littlewood(1962), eminent early 20th century mathematician, on the Riemann hypothesis: "I believe this to be false. There is no evidence at all for it .... There is no imaginable reason why it should be true."

25 Not long ago the four colour theorem/problem and Fermat’s conjecture would have generated other examples. There are two loose connections with Gödel's theorem: that there can be no comprehensive mechanical way of telling whether something is mathematically provable, and that inductive evidence of a system’s consistency can coexist with fundamental obstacles to proving it.

of the relevant systems and on that basis have known that it was true, while being unaware of Noether's 1915 theorem showing this to be a mathematical fact. Various propositions for which there is empirical evidence, including evolution by natural selection, have been claimed to be mathematically provable, and, whether or not they really are, the claim is intelligible. And, just as there is overwhelming inductive evidence that with two apples and three oranges you have five fruit, someone could have accumulated enough evidence to accept that when you walk up a hill and then down it again there is always a point which you reach at exactly the same time on both journeys, without realizing that empirical evidence was not really necessary as it could have been proven.

(C) Then there are claims that are at one time accepted as part of mathematics but are later thought to be truths about the physical world. In ancient times arithmetic and geometry were considered to be very similar; truths in either were to be be established by proof alone, as were many connections between them. Suppose that space actually is Euclidian. Do we need proofs for its properties when we can verify them with multiple observations?

(D) A hard to classify case is that of the consistency of powerful set theories. We often have considerable evidence that a set theory such as the standard ZFC is consistent, from our having worked with it and encountering no contradictions, but we also know that we will not be able to produce a proof of its consistency using resources that are no less plausible than that set theory itself.[26] So in a clear but easily misunderstood way the consistency is unprovable. Again we have something that it is reasonable to believe and which we expect to be true in a wide range of possible worlds. But some will balk at saying that we know it.[27]

We make and react differently to claims of all these kinds depending on whether we think we know them and whether we think the content is mathematical. We hesitate and hedge

---

26 Why set theory rather than arithmetic, also incompletely axiomatized? Because of the different position of informal proof. It is somewhat unclear that we cannot know by informal mathematical means that, say, first order Peano arithmetic is consistent, while the further reaches of set theory are so unintuitive that we have to depend on carefully formalized proof.

27 I am not claiming to be one of them. The question of whether belief in very powerful systems of set theory, when they are consistent, is knowledge involves delicate issues about degrees of necessity.

and tread carefully around claims about knowledge when we take the content to be mathematical but lack anything with the solidity of a proof meeting contemporary standards. Why is this?

An answer lies in the uniform-but-variable account just outlined. When we describe a person's conclusion as knowledge of some topic, we connect the description with other instances of knowledge within a framework that takes account of what is possible for us, what we can hope to achieve, on that topic. So when we characterize the topic as mathematical we suggest that it has features that trigger certain criteria. And of course we might be wrong about this. More specifically, mathematical knowledge, as described here, requires that the exceptions to perfect correlation concentrate on cases where a belief obtained with the same method as the one in question is false (rather than cases where a situation like the one that obtained is not accompanied by a matching belief). Knowledge is always knowledge, but it is a human thing subject to our realistic aspirations. Importantly, the triggers for the fine-tuning of knowledge are independent of what label we apply to a belief's supposed .domain. They depend on the facts about how one can get a reliably true belief about the entities and attributes in question. As a result, we may think that one form of fine-tuning is appropriate while in fact another is.

In (A) cases, where conjectures about numbers are settled for many instances, the method proposed, simple induction, will of course also produce false conclusions, and many of these will be much less necessary than the actual fact would be. So the classification that would fit inductive generalizations will not fit the possibilities of the case. With (B) cases, too, the claim and its support do not mesh in terms of necessity, resulting in the same tension between the epistemic classification and its metaphysical ground.

(C) cases are in many ways the opposite. We learn — on empirical grounds — that parts of the topic are less necessary than we had thought. The original methods continue to be applicable (though other, more empirical, methods may be added to them). As long as physical space is taken as Euclidean, establishing its properties by the old synthetic or analytic methods remains an option. After all, if these methods continue to be reliable sources of conclusions that combine well with those we learn about in other ways it would be foolish to spurn them. So to some extent the demotion is just from a category of necessities that can be established only by mathematical methods to a category of necessities that also allows more standard scientific evidence. What is happening could be described as a redrawing of the boundaries of mathematical knowledge, in accordance with a stratification of necessity. Facts about numbers and structures of number-like entitles are put into a category where they would be true even if geometrical facts failed. But *if* geometry is still taken as true of the world, we have a thinking-based source of conclusions which we can still take to give us truths of some degree of necessity.

(D) cases are tricky. Our confidence that a set of axioms for set theory is consistent does depend on properly mathematical practice. But it is based, in a generally inductive way, on the fact that we have arrived at a large but finite set of results rather than by *going* further by continuing from them. As a result for every such argument there is a parallel very similar argument to the consistency of an inconsistent variant of the axioms. So though mathematical in content the belief in the theory's consistency is knowledge, when it is, in a way different from that of the theory's consequences, when they are.

Mathematical knowledge emerges as distinct from, but unified with, other kinds. The label of mathematical does not play a role in getting to this conclusion. What is important is that mathematical conclusions are based on characteristically mathematical activities which give results of a characteristic necessity. So our ascriptions depend on what we take these activities and the associated necessity to be. As long as they match they will constitute a stable concept of knowledge, but given a particular time and topic communicative convenience will force a particular combination into association with the topic. It makes one wonder what future reallocations of method and modal category there may be.

***doing* math**

What are the actions whose success or failure is a basis for knowledge in mathematics? They are very varied. It is a theme of recent philosophy of mathematics that mathematical activity consists of a lot more than proving theorems from axioms.[28]

28 Azzouni (2006), part II.

Constructing examples, refining intuitions, seeing analogies between one structure and another, and many other processes are central. Unless a particular agent has performed suitable actions — which ones depending on the person and the conclusion — her thinking may not take a form where its near variants lead to truths. As many writers have pointed out,proof itself is a rather fluid business and is not in stark contrast with intuition.[29] The next step in a proof typically relies on assumptions which the practising mathematician may not be able to state, and it typically requires some training and aptitude to be able to make this step and recognize it as acceptable.[30] (To put it differently, automated theorem proving doesn’t cover the same ground, and its output is often practically unintelligible.[31])

A particularly telling aspect of mathematical practice is construction. There is a long history in geometry of producing or transforming figures in various specified ways with various means, which are used to establish associated results. Plane constructions using procedures identified with ruler and compass alone, which have sophisticated algebraic connections, are typical. What one can make and what one can prove are closely connected here.

Some of these practices are social. They can be explicit interactions: mathematicians suggest problems to others, and criticize, refine, and reformulate one another's work. One person may have a fertile imagination and suggest possibilities together with rough and often unworkable ideas about how to prove them, and then one or more other people can massage them into forms that will withstand the scrutiny of yet more people. It may be that only in the context of this network are the surviving actual and idea-sketches of the first person typically true. The available vocabulary and notation may make a fundamental difference to how one person can fix in their mind what might be worth making more precise or finding a proof for and in consulting with others for this. It may be that without some detail of procedure or notation it would have been impossible

29 For example Azzouni (2004), Larvor (2012).

30 In this connection note the suggestion, in van der Waerden (1954, chapter VIII) that ancient geometry required a learned use of diagrams which when lost prevented the continuation of the tradition. Netz (1999, especially chapter 1) presents a congruent, but more careful, position. See also Azzouni (2004).

31 There are formally correct machine proofs that humans cannot follow. But a difference between computer proof and computer chess is that recognising a checkmate is trivial while seeing that a conclusion in mathematics-as-practiced is correct can be far from it.

for some individual or for individuals in some context to keep a safe distance between their actual production and possible false output. One obvious aspect here is the process of successive reformulation by which conjectures and strategies for proving them generate reformulations and more solid proofs until there is eventually a close fit between what is claimed and what is demonstrated. A capacity to come up with suggestions that are amenable to proof or counterexample is needed, as are capacities of finding the proofs and of fitting revised proofs to revised suggestions. There are no basic differences in this respect between knowledge in mathematics and in other kinds of knowledge.[32]

Different thinkers working in different traditions in different intellectual/social surroundings need different ways of making their thinking reliable. We should not be quick to deny that Euclid knew that the interior angles triangle sum to 180°, although his axioms were deficient and he mistakenly thought they applied to the space of familiar experience. We should not be quick to deny that Leibniz knew the product rule for derivatives, although he had nothing like what we would consider a proof. We should not be quick to deny that Cantor knew that there are more real numbers than integers, although he was reasoning informally (but rigorously) from principles that we can now see to be contradictory. We should not be quick to deny that Ramunajan knew that his infinite series converged to π, although he would have disdained a conventional proof. There are many many examples.[33] In all of these cases what we would now call proof is lacking. These mathematicians might well have balked at giving more tedious proofs, and in any case they would not have known the rules as we now give them. But they and their contemporaries, partly because of their ways of interacting, knew a lot.

---

32 There is an obvious link here with the neglected topic of mathematical knowledge by testimony. But conjectures, suggestions, criticisms, and analogies are also transmitted. These interact in hard-to-understand ways.

33 One class of historically important cases concern results that need to be qualified with conditions distinguishing between kinds of convergence and kinds of continuity, that were not in the mathematical atmosphere until the the nineteenth century, in part because stating the conditions requires meticulous use of quantifiers that would have been alien until then. (Without which consistent axioms for set theory would be impossible: a hidden methodological continuity.) The work of mathematicians working in the days before careful foundations of analysis and before formal measure theory needed some retrospective tidying up and some results needed to be restated with qualifications, but the overwhelming majority was uncontroversial knowledge, for all that. A nice discussion of these issues, showing the relevance of the formal issues to applications, is in Bressoud (2007). See also Gray (2015).

causation revisited

The account has been causal, in that it has appealed to causal processes in the minds of individual thinkers and social cognitive interactions between thinkers, leading to knowledge of facts that are beyond causality. That does not make direct causal connections between mathematical facts and mathematical belief. But there are indirect connections, which are in a definite way causal. These indirect connections can be summed up by saying that mathematical facts *shape* or *constrain* processes that produce mathematical beliefs; it is because these processes respect the relevant facts that they result in corresponding true beliefs. This is still somewhat metaphorical, but it can be sharpened.

David Lewis taught us that conditionals of the form "if *cause* hadn't happened then *effect* wouldn't have followed" are particularly relevant to issues of causation.[34] But these are directed at individual physical events occurring at particular times. And "if 7+4 were not 11" is a pretty indigestible antecedent. Yet necessary general truths have contingent particular consequences, often closely linked to them. Geometrical examples are typical and abundant. If we had not travelled half the circumference of the earth on each leg of the trip we would not have traversed a triangle with three 90° angles. Everyday oranges and apples arithmetic is also full of them. If there had not been five dogs as well as the seven cats on the list of veterinary appointments then the list would not have had twelve animals, given that the vet was only seeing dogs and cats that day. But it is not necessary that there were five dogs and seven cats. Background conditions are generally like this. Universal gravitation is causally relevant to the orbit of a planet, although conditionals beginning "if there were no gravitational attraction between all massive bodies then …" are hard to handle. We think instead in terms of conditionals saying such things as that if the relative masses, positions, and velocities of a particular planet and its star had not been as they were then the orbit would not have been as it was. The causal relevance of general facts, many of them laws of nature, is no more puzzling than the rest of causation, but it is different from the causal relation between events. The relevance is focused: a law can give the reason why some events occur but not others.[35]

34 Lewis (1973), (2004). Lewis begins with this contrapositive conditional and extends it in various ways, different at different stages of his thinking on the topic.

35 I am using the vocabulary of laws for convenience, but not much hangs on what form we take general natural necessities to have.

Electromagnetic theory, though in the same nomological category as gravitation, does not generate counterfactuals with the same relevance to planetary orbits. And in general causal necessities have contingent consequences which allow them to shape individual events.

The word "causal" can be misleading here. It is more general than "causes", taken to denote a relation between token events and narrow classes of them. For it also applies to the relation between laws of nature and the events that conform to them. In fact, there can be laws that are thoroughly causal but do not give causes[36]. I shall describe theories as causal when they are true or false because of the physical workings of the world, and keep "cause" and "causes" for the relation between events. Causal theories include accounts of background conditions for event causation, that is, as I shall use the term, facts about laws of nature and similar determinants of what happens which do not satisfy the Lewisian "if not then not" formula..

Mathematical facts are background conditions for the causal relations between the stages of a person's thinking, and are invoked in causal theories of them. Only some consequences of some particular mathematical facts will be relevant to establishing a particular contingent counterfactual. Even taking mathematical propositions or facts as indistinguishable from one another, all holding in all possible worlds, most of calculus, geometry, and topology is relevant to thinking about the vet's appointments only inasmuch as they entail basic arithmetic. On the other hand basic arithmetic applies directly. (Other parts of mathematics have to reach through arithmetic to connect with everyday 'how many' thinking.[37]) And similarly in other cases: only instances of some specific mathematical principles are relevant to any particular causal application. When the topic is the epistemology of mathematics, the important question is how the thinking that leads to a particular conclusion is shaped by the particular fact that the conclusion corresponds to.

---

36 A suggestion of Mark Wilson's (Wilson 2017) is relevant here. Wilson points out that physical processes which fit the image of causes prompting their effects in time tend to be expressed by hyperbolic differential equations, while other equally fundamental natural principles are given by elliptic differential equations. So, crudely, if "causal" refers to the workings of nature in general, then the transition from cause to effect is only a very particular part of the operation of the world.

37 "Arithmetic" here means any way of doing arithmetic operations, not specifically axiomatic system and including arithmetic embedded in first-order logic with identity.

This is easy to see in some particular cases. In the mind of someone doing basic arithmetic with a mental abacus, images of unit tokens are moved around to form sums and products. This is like dogs and cats at the vets office in the example above, where possible arrangements are constrained by directly corresponding arithmetical facts. And extending this to the kind of mental rearrangement that is implicit in procedures like those that little Gauss must have gone through to find a shortcut for seeing that the sum of the integers from 1 to 100 must be half of 100×101 would be very natural.[38] Something similar will be true of any broadly analog thinking, such as simple Euclidean geometry.[39] There is an obvious connection with some kinds of mathematical intuition.

There are two general connections with mathematical knowledge. The first is sensitivity to counterexamples. Someone notices that the digits of the first several multiples of nine when expressed in the usual base ten notation add up to nine. She wonders if this is always true, and tries first to check for counterexamples. The digits of 9×10 sum to a number whose digits sum to 9, leading to the conjecture that repeated such summations will eventually lead to 9. After checking a number of these she tries to prove that it is true in general. A little algebra convinces her that it is. If all has gone well she then knows that it is. But her conclusion would not be knowledge if amid her consideration of possible counterexamples there was an easily accessed case that she would, or even easily could have, mistakenly taken to be a counterexample. Or if she would have been unmoved by counterexamples to a very similar but invalid generalization. And part of the reason why she will not come across counterexamples if she is on the way to knowing her conclusion is that in fact the sum of sums *is* always 9.[40]

Sometimes the way that the truth of a conclusion shapes the reasoning leading to it is transparent. An example is a proof of 2+ 1 = 3 in Peano arithmetic. One line in a simple

38 From a famous story that when Gauss was a schoolchild his teacher, hoping for a quiet hour, set the class to add the integers from 1 to 100, not anticipating that within minutes the boy would have realized that if you put the sequence beside its reverse and then add corresponding terms you get the same subtotal each time, so that you need only to multiply the number of terms by this constant subtotal and then divide by two.

39 Along the walls of my classroom in the first school I attended there were pictures of playing card faces in the standard patterns, and even now, many decades later, when uncertain about sums and differences, instead of remembered tables I put images of the card faces together and count the markings.

40 This is not to deny that many other mathematical facts would also play a role in explaining how a person did or did not think and how they reacted to a particular example.

proof might run "SS0+S0=SSS0", followed by definitions of the three numerals and of addition. But note how there are two occurrences of "S" and then one more occurrence on the left of the equals sign and three on the right. Had we been proving a different identity there would have been different numbers of "S"s at this stage, and had the person producing the proof have mistakenly written down the wrong number of "S"s the result would not have been a correct proof. (Proofs in simple arithmetic and logic can seem circular for this reason. Someone who did not understand the conclusion would find it hard to follow the reasoning. It is not in fact the reasoning that is circular, but the grounds for its success.) When a proof can be conveyed by a diagram, without words, there is usually a similar clear connection between the mathematics itself and making or understanding a proof of it.[41]

Most proofs are not as autonomous as this. They build on other results and on explicit axioms. Proofs of any sophistication build on other proofs and on acquaintance with suitable axioms. Knowledge transmitted from mathematician to mathematician will inevitably play a large role, with the conclusion developing at each stage. (Like a wild rumour, except for the little matter of truth.) The fact still applies. The result is a necessary condition for the success of its its proof.[42] This is best put in terms of models. A previous result on which a new proof is based will establish that its conclusion holds in a certain model, and a correct proof will preserve this.[43] If the conclusion of a proof does not hold in some model satisfying all its starting points then something has gone wrong with the reasoning. So its holding in the intended model is required for the reasoning to be good. And this fact about the abstract structures is a necessary condition for the connection between their representations. If the structures were not so related then the reasoning would not have been possible, on the assumption that it is causally guided in such a way that it preserves truth in a model. Or, to put it differently, when the structural facts do not hold the reasoning can only get to its destination when it is flawed.

---

41 Nelson (1993) and successor volumes are a rich and varied source of purely visual proofs.

42 This raises a delicate parallel to knowledge of other topics. Suppose that someone makes a correct proof of a true result based on a previous true result itself resting on a faulty proof. When is this knowledge?

43 Sometimes the intended domain will need to be expanded to accommodate a name associated with a syntactically consistent description. For example one might build a nonstandard model for arithmetic around the set of descriptions "number greater than n" for all standard numerals n. But then, everything else going well, the proof will establish something about the nonstandard model.

There is another connection with real life non-mechanical proofs. Many steps of them make more implicit assumptions than are stated, or even that could be explicitly formulated by their authors (Larvor 2012, Azzouni 2013). To that extent training or intuition — a feel for the material — is essential. And these transitions are possible only because of the facts that they unveil.

causal theories of truth and reference

I have been using "true" as an undefined primitive. That's alright; you cannot answer every follow-up question. But the capacity to refer to mathematical objects may seem just as problematic as the capacity to have knowledge of them. So there is a point to explaining, briefly, why the attitude to causation I have been outlining, centred on the role of background facts in laws describing the origins of human actions, can be turned with remarkably few changes into a version of a causal theory of truth.

The first target has to be reference. Once the initial links between names, demonstratives, and predicates and their referents are determined, reference can be passed on from one use and one speaker to another, and truth conditions can follow the same recursive patterns, given a somewhat arbitrary choice of logical form, with the same strategies and meeting the same problems for sentences about any topic. The problem is about establishing reference in the first place, rather than truth. Initial reference to a physical object links its physical characteristics to the act of naming. I point at a baby and say "this child will be called Maryam", and the infant in line with my ostention gets the name for the moment and keeps it if others take it up and pass it along. If another infant had been there and had caused my attention and that of the audience to be directed towards it then the name would have latched on to her or him instead. The puzzle is how one can name abstract objects in a mathematical context in a parallel way.

A person's attention is directed to a physical object — they have it in mind — when the fact that the object has some relation to the person, often in terms of their relative locations, is a cause of the person’s coming have a capacity to attribute properties accurately to it.[44] By naming the object and drawing another person's attention to it and to the name the connection can be passed on to another person, who can pass it on to others. The initial connection can use a description that is in fact false, as long as it leads to a capacity for truthful ascription, which can persist without the description, as in Kripke's Gödel example.

The situation with names in mathematics is parallel in a way that should by now be familiar.[45] Realizing that when the larger of two real numbers has the same ratio to the smaller as their sum has to the larger this ratio is the same whatever the numbers, someone names the ratio "φ". The relation between namer and named is that the person has identified the constant. Then others can use it as a name for its value (1.6180339887498948420 …), without necessarily knowing that it has this originally defining characteristic. (A subsequent user of "φ" may have learned it as the limit of the ratio of terms of a Fibonacci sequence, for example.) It is even possible for the initial description to be wrong, as would happen if someone introduced π as 22/7 thinking that this was the ratio of the circumference and diameter of a circle, or if someone introduced "√2" as the unique square root of 2, not realizing that -√2 also qualifies.

The characteristics of φ are causally related to the person who introduces the name "φ" . Suppose that the realization that there is such a constant comes from some particular lengths, perhaps in art or design. If the objects with those lengths had been different sizes then they would not have had a constant ratio, and whatever ratios they had would not be φ. When someone calculates the ratio of two lengths that do qualify the answer they get if they make no mistakes is identical, up to some number of decimals, with φ, and the reason is the mathematical fact that there is such a constant and its value is φ . And so on. The facts about φ are conditioning the causal possibilities of thinking about it.

Since transmission need not preserve any accurate characterization that facilitated the initial naming, a name can be used to say false things, even very false things, about its

44 This way of putting it has a loose resemblance to Jane Heal's concept of co-cognition in understanding another person's thinking (Heal 1998). The resemblance may go deep in that shared attention is sometimes associated with language learning.

45 The picture is complicated by the fact that in mathematical practice one frequently introduces a symbol which is used several times and then not again with the same denotation. (We sometimes improvise something similar in everyday language.) I am ignoring this.

referent, here as in the case of reference to physical objects. Just as someone can intelligbly say of whales that they are fish, they can say of φ that it is rational. The ascription is perfectly intelligible, but false.[46]

Such causal accounts of proper mathematical names can be easily extended to predicates. (There are relatively few proper names in mathematics, mostly names of particular constants and functions.[47] Most of the work is done by descriptions of functions.) The pattern is that a predicate or function is introduced in connection with a relationship underlying a pattern, and then continues to be used for what accounts for this pattern even when the description or functional role that was originally part of the introduction turns out to be inadequate. Exponentiation gives an elementary example. m^n originally meant the n-termed product $m_{(1)} \times \ldots \times m_{(n)}$, and although this has to be abandoned when n is not a positive integer, as with $e^{\pi i}$, the notation remains because it is convenient and because there is a unity to the mathematical origins of the numerical patterns that result. The general phenomenon, often resulting in "analytic continuation", is the extension of well-behaved functions to new domains, sometimes by using a convergent series as a bridge. This can seem at once like redefinition and like uncovering what a function really was all along.[48]

Conclusion: This has not been an argument for mathematical Platonism. Or against it. The Plato/Benacerraf problem may indeed be daunting; there may still might be a tremendous challenge in reconciling a plausible ontology for mathematics, what in the world different parts of mathematics are about, with a plausible account of the grounds for mathematical beliefs, how we can know when claims about these specific objects are true. Or there may not be. There is no inevitable tension between extending accounts of truth-in-general to mathematical language and extending accounts of knowledge-in-general to mathematical results and discoveries. Deep complications enter, if they do,

46 The greatest importance of this separation is to make it clear how one can say extremely false things about something while still talking about it The same could be said about names in fiction. And in fact the account would fit comfortably with the interpretation of fictional names in Kripke (2011) as names of abstract objects.

47 "But there are infinitely many names of integers." No:they are always and inevitably descriptions recursively generated from a few primitives.

48 Generalizing from the discrete to the continuous is a feature in many domains. Startling cases are geometries of fractional dimension and fractional derivatives and integrals.

when we go further and combine accounts of specifically mathematical ontology with accounts of specifically mathematical method. But whatever the size of this challenge, it does not make a problem reckoning mathematical lore as of a piece with other knowledge, in part because its truth is of a piece with other truth. The causal elements of both can even be transferred, if they are described in enough generality.

"Is of a piece with", "enough generality": the resulting theories are neutral on many topics that philosophers care about. That is an advantage as well as a liability. It makes us ask what purposes we want an account of a topic to serve, and emphasizes there is typically no such thing as *the* theory of X. (Consider, for example, thermodynamics and the molecular theory of heat, which ought to combine smoothly but ought also to be separable.) The bearers of truth and of knowledge, for my purposes here, have syntactical structure only as a device to get the right truth values.[49] Their structure may have little relation to that of the facts that verify or falsify them; for them the world is all that is the case.

bibliography

Azzouni, Jody (2004) Proof And Ontology in Euclidean Mathematics. In *New Trends in the History and Philosophy of Mathematics*, Kjeldsen, Pedersen, and Sonne-Hansen, eds., 117-133: Denmark: University Press of Southern Denmark.

Azzouni, Jody. (2013) The relationship of derivations in artificial languages to ordinary rigorous mathematical proof. *Philosophia Mathematica* (**III**) 21: 247–254.

Benacerraf, Paul (1973) Mathematical Truth. *The Journal of Philosophy*, **70**, 19: 661-679

Bennett, Jonathan (2003) *A Philosophical Guide to Conditionals*. Oxford: Clarendon Press.

Bressoud, David (2007) *A Radical Approach to Real Analysis.* The Mathematical Association of America.

Buijsman, Stefan (2019) Learning the Natural Numbers as a Child. *Noûs* **53**, 1: 3–22

49 A snappy way of putting this is that propositions are reconstructed to make sense of what we do with them. Logical form is just a device to make this tractable. Compare Wilson (1992).

Callard, Benjamin (2018) Can Math Move Matter? *Inquiry: An Interdisciplinary Journal of Philosophy*:1-26 (forthcoming)

Colyvan, Mark (2019) Indispensability Arguments in the Philosophy of Mathematics, *The Stanford Encyclopedia of Philosophy*(Spring 2019 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2019/entries/mathphil-indis/>.

Craig, William (1990) *Knowledge and the state of nature*. Oxford: Oxford University Press

Dehaene, Stanislas, E. Spelke, P. Pinel, R. Stanescu and S. Tsivkin(1999) Sources of Mathematical Thinking: Behavioral and Brain-Imaging Evidence. S*cience*, New Series, **284**, 5416: 970-974.

Dehaene Stanislas (2009) Origins of Mathematical Intuitions. *The Year in Cognitive Neuroscience Ann. N.Y. Acad. Sci*. 1156: 232–259.

Dehaene, Stanislas (2004) Evolution of human cortical circuits for reading and arithmetic. In S. Dehaene, J. R. Duhamel, M. Hauser & G. Rizzolatti (Eds.), *From monkey brain to human brain*. Cambridge, Massachusetts: MIT Press.

Ferreirós, José Mathematical Knowledge and Practices. In M. Suarez et al. (eds.), *Philosophical Issues in the Sciences*, ch 6: 55-64.

Goldman, Alvin (1976) "Discrimination and perceptual knowledge", *Journal of Philosophy* 73: 771-91.

Greco, Daniel (2016) Safety, Explanation, Iteration. *Philosophical Issues* **26**, 1:187-208

Gupta, Anil (1993) *A Critique of Deflationism*. *Philosophical Topics*, **21**, 2: 57-81

Hale, Robert and Crispin Wright (2002) Benacerraf's Dilemma Revisited. *The European Journal of Philosophy* **10**, 1: 101-129.

Heal, Jane (1998) Co-cognition and off-line simulation*. Mind and Language* **13**, 4: 348-64.

Jubien, Michael (1977) Ontology and Mathematical Truth. *Noûs*, **11**, 2: 133-150.

Kanamori (2018) Aspect-Perception and the History of Mathematics. In Michael Beaney, Brendan Harrington, Dominic Shaw, eds. *Aspect Perception after Wittgenstein: Seeing-As and Novelty*: 109-130.

Kitcher, Philip (1980) A Priori Knowledge. *The Philosophical Review*, **89**, 1: 3-23.

Kitcher, Philip (1988) Mathematical Naturalism. In *Essays on the History and Philosophy of Modern Mathematics*, W. Aspray and P. Kitcher eds. Minnesota Studies in the Philosophy of Science 11: 293–328.

Larvor, Brendan (2012) How to think about informal proofs. *Synthese*, **187**, 2: 715-730.

Lewis , David (1973) *Counterfactuals*. Cambridge, Mass.: Harvard University Press.

Maddy, Penelope (1980) Perception and Mathematical Intuition. *The Philosophical Review*, **89**, 2: 163-196.

Margolis, Eric (2019) The Small Number System. *Philosophy of Science,* forthcoming.

Netz, Reviel (1999) *The Shaping of Deduction in Greek Mathematics*. Cambridge: Cambridge University Press.

Nelsen, Roger B. (1993) *Proofs without Words: Exercises in Visual Thinking.* Mathematical Association of America.

Pritchard, Duncan (2015) Anti-Luck Epistemology and the Gettier Problem. *Philosophical Studies* **17**, 1: 93-111.

Quine, W. V. (1981) Success and Limits of Mathematization" in *Theories and Things*, Cambridge, MA: Harvard University Press: 148–155.

Rabinowitz (web) The Safety Condition for Knowledge. *Internet Encyclopedia of Philosophy* (https://www.iep.utm.edu/safety-c/)

Resnik, Michael (1975) Mathematical Knowledge and Pattern Cognition. *Canadian Journal of Philosophy*, **5,,** 1: 25-39.

Sosa, Ernest (1999) How must knowledge be related to what is known? *Philosophical topics* **26**, 1 &2: 373-384.

Soylu, F., Lester, F.K., & Newman, S.D. (2018) You can count on your fingers: The role of fingers in early mathematical development. *Journal of Numerical Cognition*, **4**, 1: 107–135.

Tait, W. W. (1968) Truth and Proof: the Platonism of Mathematics. *Synthese* **69**, 3: 341-370.

Tappenden, James (1995) Extending Knowledge and `Fruitful Concepts':Fregean Themes in the Foundations of Mathematics. *Noûs* **29**, 4: 427-467.

van der Waerden, B L (1954) *Science Awakening.* Noordhoff: Groningen.

Welbourne, Michael (1986) *The Community of Knowledge.* Aberdeen: Aberdeen University Press.

Wilson Mark (1994) Can We Trust Logical Form? *The Journal of Philosophy*, **91**: 519-544.

Wilson, Mark (2017) *Physics Avoidance*. Oxford: Oxford University Press.