
THE REASONER

VOLUME 15, NUMBER 4
JULY-AUGUST 2021

thereasoner.org
ISSN 1757-0522


CONTENTS

Editorial	27
Features	27
News	31
What's Hot in ...	31
Events	33
Courses and Programmes	33
Jobs and Studentships	34

EDITORIAL

Dear Reasoners,

In the interview which opens this issue, Kevin Knight has been kind enough to share his thoughts on taking up a non-academic career, among many other things. When I started my PhD in 2002, Kevin was the senior occupant of the Mathematical Logic PhD office, on the fifteenth floor of the Manchester Mathematics Tower. (Just in case you are tempted to look it up, that Oxford Street landmark is no longer there, and the Department has been relocated to the much less architecturally controversial Alan Turing Building.) Like many other Logic PhD graduates of my acquaintance, Kevin established



his career in industry, where his background turned out to be quite relevant, as you can see below. I hope that his experience may provide useful inspiration to those readers who are currently making up their mind on what to do next.

HYKEL HOSNI
University of Milan

FEATURES

Interview with Kevin Knight	33
HYKEL HOSNI: Can you tell us something about your background?	33
KEVIN KNIGHT: I did my undergraduate degree at Rice University, where I had three majors: Mathematics, Computational & Applied Mathematics, and Philosophy. Obviously, there was a lot of overlap between the first two. In mathematics, my interests were at sort of opposite ends of the spectrum: numerical methods and logic. In philosophy, my interests were also sort of at opposite ends of the spectrum: political philosophy and logic. It doesn't take a logician to see the intersection of those two sets.	34

HH: It's definitely non empty! So you went for it ...

KK: No! I didn't go straight from Rice to graduate school in logic. At some point I learned that I could teach political and legal philosophy as either a philosophy professor or a law professor. In the US, at least, law professors make significantly more money than philosophy professors. I reasoned that if I had an option to get more money for doing the same thing, I should take that option.

HH: Sure-thing reasoning – a logical enclave in economics territory.

KK: Almost. I enrolled in a joint law and philosophy program (JD/MA) at New York University. The program started with a year of law school, and I realized very quickly that I did not want to be there.

HH: The club of logicians who moved away from law is well populated, at least since the times of Leibniz.

KK: Indeed, I realized soon enough to apply for graduate schools in mathematics and philosophy for the next academic year.

HH: How did you choose then?

KK: I started by looking at the people who had written my undergraduate math textbooks. As it turns out, I had a numerical analysis textbook written by Nicholas Higham, a professor at the University of Manchester. When looking at the Department there, I saw that they had a logic group. So I applied. It's kind of funny that it was my interest in applied mathematics that led me to Manchester's program in logic.

HH: There is probably a common cause by the name of Alan Mathison Turing behind this apparent coincidence. Did you just apply to Manchester?

KK: I no longer remember the full set of programs I applied to, but I recall that I was accepted at three: Manchester (mathematics), Leeds (mathematics), and the University of Texas at Austin (philosophy). At that point, I'd lived most of my life in Texas, so I decided I would choose one of the programs in the UK.

HH: Fair enough, did you toss a coin then?

KK: No, I chose to study with Jeff Paris at Manchester over the program at Leeds for a few reasons. First, when I asked my logic professor from Rice (Richard Grandy) what he thought about the options, Jeff was the only person whose work he knew anything about. Second, my interest in logic was more on the practical side, and Jeff's research interests seemed more in line with that than what was going on at Leeds.

HH: What do you mean by "the practical side" of logic? Your Rice professor probably was acquainted with the undecidability of Ramsey Theorem.

KK: What I mean is that my main interest was in how reasoning can/should be done. I've never really been interested in the more theoretical parts, like model theory or set theory. I've also been principally interested in reasoning where the proofs are human understandable - at least where humans can understand all the individual steps in the proofs, even if a proof is too long for a single human to fully understand. That's why I always restricted myself to proof systems where the rules were simple.

HH: This is practical logic as some would say. By why did you specialise in *paraconsistent logic* then?

KK: One reason is that it was not a very crowded field. I didn't have to worry about a lot of people working on the same problems I was working on and maybe getting answers first. Another reason is that I like paradoxes. I don't know why, exactly, but they seem fun to me. I especially like the kind where a set of highly plausible premises leads to a contradiction. Also, it seemed like a practical problem.

HH: In the sense that some contradictions are true?

KK: Well, we do encounter inconsistencies all the time in real life. In the end, though, I think my work in paraconsistent logic wasn't all that practical. Few, if any, inconsistencies that you find in practice are of the variety that you find in philosophical paradoxes. Most are either data entry errors (in which case you want to correct the error) or different sources making conflicting claims (in which case you want to contextualize the information by source). In neither of those cases, do you really want to compensate by adjusting your rules of logic.

HH: That's a very classical-logical analysis.

KK: Still, I enjoyed the subject matter of my PhD, and I learned a lot. I learned, for example, how to decompose my work into manageable chunks. In the case of my thesis, that was a matter of breaking things down into lemmas and propositions, which could be individually understood and potentially re-used, rather than just going straight to the big theorems. That helped me greatly in my career as a programmer, as decomposing your functionality into useful units is key to writing good code.

HH: This is very interesting. After your PhD you went back to the US and took a job in industry. Before telling us about that, I guess our readers, especially PhD students, would be very interested in knowing whether any specific thing which you learnt during your PhD turned out to be particularly useful in your career.

KK: Yes. One thing that I learned in particular from Jeff was how to approach modeling something. In his case, he was trying to model uncertain reasoning. He started by setting out the properties he thought his model should have, and then proved that his model of uncertain reasoning satisfied these properties. I think he got this approach from Claude Shannon, who used a similar approach for his information measure. You can generalize this idea by saying that you should start modeling by thinking about what you want out of the model and how it should behave. And then you make sure that the model that you come up with meets these expectations.

HH: I guess this bridges the gap between the abstraction of logic and its practical applicability: you need to know what you where you want to get before setting out to find how to get there. The perception of mathematical models outside the academia is not always of this kind, it seems to me.

KK: Over my career, I've seen a number of people not approach modeling this way, and they often end up with models that don't behave in sensible ways or otherwise aren't usable.

HH: Speaking of your career, was it always clear to you that you it would be in industry? KK: No. I actually had intended to go into academics. But I didn't start looking for jobs until I finished my PhD, which was in July. As it turns out, there aren't many academic jobs available in July. The positions that are available are the ones that they couldn't fill during the normal hiring season, usually because they weren't very desirable. Still, I applied to a few (I forget where), but I never heard back on any of them. I doubt they were very interested in a specialist in mathematical logic, much less paraconsistent logic. But I needed a job, and I couldn't wait around for the next academic year. So I started applying for jobs at various businesses. Again, I can't remember all the places, but I didn't hear back from most of them.

HH: But then you got hired in Austin, Texas by [Cycorp](#) – an AI legend.

KK: Yes. When I joined, I think there were 60-70 employees in total. I was at the company for a little over 4 years. They did logic based artificial intelligence with a system called Cyc. The aim was to build an AI with a large common sense knowledge base. That is, broad knowledge about how the world



works in an everyday sense, rather than deep and precise scientific knowledge. For example, knowing the fact that an electric kettle will boil water and the steps to make it happen (plug it in, put in water, turn it on), rather than knowing the physics of how it does it. A good portion of the employees had advanced degrees in philosophy and were tasked with modeling this kind of knowledge.

HH: The idea underlying the Cyc project is very interesting to look at in the era of AlphaFold, GPT-3 and convolutional Neural Networks. At Cyc there hired humans doing the models!

KK: I did a certain amount of this kind of modeling, but mostly I worked on Cyc's inference engine. Cyc had a general inference engine coupled with a number of specialized modules for particular types of reasoning. I worked on, among other things, temporal and subsumption-based reasoning modules. I also worked on the module that allowed Cyc to use external sources to get information (such as a database of geographical information or a weather reporting website). The platform was written in a dialect of LISP, and Cyc had its own modeling language called CycL (which was basically a language of Nth order predicate logic).

HH: Since you've been there, I'd like to press you a bit more on the Cyc interpretation of AI versus the new sub-symbolic spring.

KK: The state of machine learning at the time I was at Cycorp was very different from today. My limited experience with machine learning at the time left me with the impression that it was expensive and largely impractical. But the computing landscape today is vastly different from what it was back then, and a lot of things that weren't practical then are now.

HH: GPUs and *lots* of data . . .

KK: Yes. Still, I think there's a place for human modelers and, in particular, high quality, well curated models. I'm not just talking about Cyc here, I mean this more generally. For example, one class of applications is when two or more parties need to communicate about some subject in a structured way (where the messages are programmatically understandable). The different parties need to have an agreed model for this structured communication. And in some cases, the models may even need to conform to certain laws or regulations. In cases like that, you definitely want humans creating, editing, and reviewing the models. And you want to think carefully about the model from the start, because changing the model later on can be costly (particularly if the changes are not backward compatible).

HH: This kind of modelling is not done once and for all, it's best thought of as a process

KK: Any model that's in real use will undergo changes over time. For such models, the lifecycle of development is very important: who is allowed to make changes; who is required to review and approve changes; how changes are tracked; how and when new versions of the model are released, published, and deployed. In some cases, there may even be legal or regulatory requirements around these processes. This may seem uninteresting from a more theoretical viewpoint, but having decent development processes is absolutely necessary for maintaining high quality models in the long term. Necessary, but not sufficient. You also need the people involved to take model development seriously and be disciplined about their roles in the process.

HH: After Cycorp you moved on to the University of Penn-

sylvania, in Philadelphia.

KK: Yes, in a research group in systems engineering. It was a small group. At its largest, I think there were 12 of us. I worked there for a little under 4 years. The group was made up of a professor of systems engineering, Barry Silverman, some PhD students, a few post-docs, and a handful of professional programmers (including myself). The main system I worked on there was called PMFserv (Performance Moderator Functions Server). This was a platform where social science theories could be modeled and tested using agent based simulations. In this context, an agent could be any kind of entity: a human, a state, a faction within a state, etc. So social science theories at a variety of levels could be implemented. For example, theories of individual human behavior from psychology, theories of group behavior from political science, or theories of how individual behavior affects group behavior from sociology. Part of the idea was to have a platform where people could test out different theories in the same environment to see if and how they worked together. The platform was written in Python.

HH: That sounded like the perfect tradeoff between your practical and philosophical interests. But then you moved into finance.

KK: Currently, I work at Goldman Sachs (GS), where I've been for the last 10 years. It's a very large company, with somewhere over 30,000 employees. I can't give specifics of my work at GS, except that we've recently open sourced large parts of the project that I've been working on for the last several years. It's called Legend, and it's been open sourced through an organization called FINOS: <https://www.finos.org/legend>. It comprises a suite of tools for creating conceptual models of data and performing actions based on those models.

HH: Can you give us an idea about what it does? KK: Sure, a common use case is to query data using the model as the query language so that users, especially non-technical users, do not need to worry about the particulars of the data storage. The platform was written primarily in Java, and includes its own modeling language. While the work at my previous jobs could be called research, my work at GS really is not. I work on systems that are in active use. A key point about the modeling environment is that it's designed to be accessible to people who aren't modeling specialists. That's important because most of the modelers are domain experts (people knowledgeable about the relevant business concepts or the underlying data), rather than general modeling specialists. The modeling language is much closer to UML than to predicate logic, and that turns out to be much more natural for most of our users (both in terms of creating and understanding models). I have to stress again how important usability is. It's no good having a modeling platform that people can't use.

HH: It really sounds like spending your PhD on modelling "common sense reasoning" from the logical point of view gave you a mindset, in addition to mathematical tools, which turned out to be quite useful in your career.

KK: The common thread in my career has been modeling platforms. At Cycorp, it was knowledge modeling. At UPenn, it was social science theory modeling. At GS, it has been data modeling. That wasn't something I set out to do, but it's something I noticed looking back. I enjoyed the more research oriented parts of my career, but there's also a satisfaction that comes from building a system that's in active use.

HH: Is there any explicit advice that you think could be useful to recent PhDs who are wrestling with the academic and

non-academic options for their next step?

KK: For those thinking about academic versus non-academic options, I don't have too much more to say. I guess I should reiterate the mistake I made by not starting my job search until after I finished my PhD. That cost me any hope of getting an academic position, and it probably adversely affected my ability to get other jobs as well. A number of companies, especially large ones, have certain times of year when they do a lot of their hiring. I was lucky to get the job I did at the time I did.

I would also add that it's important to realize that there is a very wide variety of non-academic options. There are many different industries. There are companies of all different sizes and cultures. There are all kinds of positions doing all kinds of different work. And you might find interesting work in unlikely places. When I got my PhD, I never would have thought I would be working in the financial industry, but I've been here for a decade now. So when searching for non-academic jobs, cast your net wide.

Dialetheism and Modus Tollens

According to dialetheists, some contradictions are true. In other words, according to dialetheists there is a sentence A such that it and its negation $\neg A$ are both true. Dialetheists have argued for many examples of true contradictions. But for the sake of illustration, we will employ a purported example of a true contradiction arising from motion, given by Graham Priest (2006: *In Contradiction*, Oxford University Press, 161), who writes: "I am in a room. As I walk through the door, am I in the room or out of (not in) it?" Priest argues that the answer to this question is a true contradiction – as I walk through the door, I'm in the room, but also not in the room.

Modus tollendo tollens, or modus tollens for short, is the argument form $A \rightarrow B, \neg B \therefore \neg A$, in which the major premise is a conditional, the minor premise the negation of its consequent, and the conclusion the negation of its antecedent. For example, the argument 'if it's raining, then it's cloudy; it's not cloudy; therefore, it's not raining' is an instance of modus tollens where the major premise is the conditional 'if it's raining, then it's cloudy', the minor premise is 'it's not cloudy', the negation of the consequent, and the conclusion is 'it's not raining', the negation of the antecedent.

We argue that if dialetheism is true, then modus tollens is invalid. To see why, suppose that both 'I'm in the room' and its negation 'I'm not in the room' are true, but that 'I'm walking through the door' is true while 'I'm not walking through the door' is not. Then consider the following instance of modus tollens:

1. If I'm walking through the door, I am in the room.
2. I am not in the room.
3. \therefore I'm not walking through the door.

By hypothesis, the second premise is true while the conclusion is not. In the scenario described, the first premise is intuitively true as well. But if so, we have an instance of modus tollens with true premises and a false conclusion, so modus tollens is invalid.

Whereas we think the instance of modus tollens described above has true premises and a false conclusion, we don't have similar intuitions about any corresponding instance of modus

ponens (the argument form $A \rightarrow B, A \therefore B$). Consider, for example, the following instance of modus ponens:

1. If I am in the room, I'm in bed.
2. I am in the room.
3. \therefore I'm in bed.

If 'I'm in the room' is both true and false, but 'I'm in bed' is just false, then 'if I am in the room, I'm in bed' is just false. Unlike the corresponding instance of modus tollens, this instance of modus ponens is intuitively valid.

The issue is notable for at least two reasons. First, many putatively invalid instances of modus tollens have been given even in the classical context. Most of these putative counterexamples involve an embedded conditional or modal in the consequent of the major premise, which is putatively negated in the minor premise (for an overview, see Theresa Helke (2018: *On Conditionals*, PhD Thesis, National University of Singapore)). But the logical complexity of the embedded conditional or modal makes it arguable whether these examples are truly instances of modus tollens. In contrast, the example we gave above is a perfectly straightforward instance of modus tollens.

Secondly, it's often observed that "one person's modus ponens is another's modus tollens". In the context of classical logic, the validity of modus ponens and modus tollens stand and fall together (as recently emphasised in this journal by Lina Lissia (2020: 'On some analogies between the counterexamples to modus ponens (and modus tollens)', *The Reasoner*, 35-7)). But we have just seen that in the context of dialetheism, the validity of modus ponens and modus tollens can come apart – some instances of modus tollens, but not of modus ponens, are intuitively invalid.

BEN BLUMSON & THERESA HELKE

Did Socrates know how to see your middle eye?

"If the inscription took our eyes to be men and advised them, 'See thyself,' how would we understand such advice?"—Socrates (apud Plato, *Alcibiades*)




Figure 1: Left: How one usually looks at oneself in a mirror. Middle: The cyclops that appears when one stands nose-to-nose with a mirror. Right: The third eye obtained by slowly backing away from the cyclops while continuing to point one's eyes at the cyclops's eye.

In an enlightening paper, Gallagher and Tsuchiya (2020: Third-Eye Rivalry, *i-Perception* 11(4), 1–8) describe a visual phenomenon in which, by properly directing one's eyes, one can perceive a clear and stable illusory third eye in one's own reflection or in the face of a colleague. Gallagher and Tsuchiya indicate some surprise at the apparent silence of the literature on this phenomenon (it is indeed surprising considering nothing would prevent its discovery thousands of years ago). We will describe this phenomenon in our own words (see

Figure 1) and then speculate that it might have been known to Socrates.

If you direct your left eye at the image of your left eye in a mirror, and your right eye at the image of your right eye, then you do literally look yourself in the eyes (plural), making each eye “see itself” (to quote Socrates’ variation on the Delphic inscription). One way to do this is to approach as close as possible to a mirror. As you approach the mirror, your eyes will tend to focus on one point (a process called *convergence*), but as you get closer and closer, it will eventually become difficult for your eyes to simultaneously track the same point. At very close proximity—when you can get no closer because of your own nose—each of your eyes will indeed look directly at itself. Just like firing an arrow at point-blank range, an eyeball so close to a desired target has little choice but to see that target!

When you stand so close to the mirror, and each eye gazes at its own reflection, the two eye-reflections will fuse, so a cyclops will stare back at you with one lone eye. This is because, in general, any time your left and right eye are directed at two similar-looking objects, those objects will fuse together in your perception—just as, looking through binoculars, you perceive one aperture, not two. Since your own two eyes look similar, when each eye sees its reflection, those reflections fuse.


And in general, whenever anything (real or illusory) is in your field of vision, you can focus your eyes on it, even while you move. Thus, when you see that illusory cyclops eye, you can consciously focus your eyes on it, while slowly stepping back from the mirror. Consciously focusing your eyes on the cyclops eye, it will persist, even as you back away. Indeed, when you focus your eyes on the cyclops eye, then you focus your left eye on your left eye’s image (for that is what the cyclops eye is to your left eye), so it will continue to point there even as you step back. Likewise, your right eye will continue to point at your right eye’s image. And as you step back, keeping the cyclops eye in focus, each eye’s field of vision will expand, until it includes the other eye’s reflection too. At that point, you will perceive three eyes: each of your eyes will perceive two eyes (a total of four), but two of those remain fused (reducing the total to three).

With practice, it is even possible to learn how to control the direction of your eyes so as to view that third eye without first coming nose-to-nose (you can even do it cross-eyed, directing each eye at the opposite eye’s image). The way we described the process, the crucial key is to direct each eye at itself: the right eye at the right eye, and the left eye simultaneously at the left. This is precisely what we would do if we were to follow the hypothetical advice in the quote at the top of this article. Socrates speaks of the Delphic inscription advising the eyes as men, plural, not as one singular man. And the advice which the inscription gives to each eye is to “see thyself” (singular). Anyone who took this advice literally (or rather, anyone whose eyes took the advice literally) would be confronted by a cyclops or by a third eye, depending on the distance to the mirror. We speculate that of all people who followed the advice literally, Socrates himself may well have been one of them.

Gallagher and Tsuchiya close their paper by quoting Shams Tabrizi, “The summary of the advice of all prophets is this; Find yourself a mirror.” If that’s true, it is abductive evidence of Socrates’ prophethood, for Diogenes Laërtius reports that Socrates “recommended to the young the constant use of the mirror, to the end that handsome men might acquire a corresponding behaviour, and ugly men conceal their defects by ed-

ucation.” We would hardly be the first in the world to suggest there are things hidden between the lines of Plato or Socrates, if we were to suggest that maybe Socrates recommended the use of mirrors for some hidden reason besides grooming, or that maybe there is more than meets the eye in Socrates’ words to Alcibiades.

SAMUEL ALEXANDER & CHRISTOPHER YANG



NEWS

Calls for Papers

CAUSAL DISCOVERY: special issue of *Transactions on Neural Networks and Learning Systems*, deadline 22 October.

WHAT’S HOT IN . . .

Science Policy

The COVID-19 pandemic challenges the whole world and requires coordinated global action. The future challenge is how to be prepared for such a global threat. One of the obvious threats of this type is global warming, while some other challenges are not easy to predict. Still, we need to incorporate preventive measures that will facilitate global action. What is specific about these global challenges is that they affect everybody and unless they are addressed equally in every part of the world, the threat will not be removed. For instance, in the case of COVID-19, even if



some countries act in a selfish way by keeping the vaccines exclusively for their citizens, there is an imminent risk of new immune evasive variants of the virus in places where it is not contained. In particular, the potential of the virus to mutate is larger the more people in any country are infected. This teaches us solidarity and that it is a priority to help the ones in the greatest need. Solidarity comes on three levels: individual, cultural, and national. Since humans are social animals who share the planet with other non-human animals, cooperative behavior seems to be a requirement for our survival. Acting, rather than reacting, requires us to be prepared for different scenarios. In addition to practicing solidarity, one of the obvious measures that will make our society more robust to global threats is the investment in science and education. As one of the big problems of the current educational system is elitism, this is something that needs to be addressed in the future. By elitism, I refer to favoring the academic systems of certain countries and to favoring specific higher education institutions. For instance, as a consequence of the exclusivity and public pressure from the ‘yellow vests’ demonstrations, the famous French *École Nationale d’Administration* will be closed. It was perceived as a prime symbol of elitism as most of its students have parents with senior roles in business or government, and the selection criteria strongly favor children with such a background. Furthermore, the curriculum, with internships at governmental departments as well as reliance on guest lectures from other elite institutions and the civil service, is reinforcing the establishment of an elite network (Nature 2021 Jun; 594(7861):7-8). During the pandemic, we witnessed that groundbreaking results in vaccine development were coming from many different countries and institutions. Moreover, we expect equal treatment and equally competent medical help in every country of the world. Thus, even if the quality of education is not everywhere on the same level, societies should strive to reach equal and high standards. Elitism is a form of manipulation in which some people feel special or chosen for no objective reason, while others are excluded from this group. On the other hand, the spirit of humanistic philosophy teaches us that everyone is able and should be encouraged to participate in the discussions related to her. The bias towards certain elitist institutions at the same time represents epistemic injustice towards the researchers and students from competing universities, institutes, and laboratories. Being open-minded towards different educational and scientific institutions is a virtue in this context. It is the basis for building epistemic trust that researchers from every part of the world are capable and responsible when it comes to finding solutions to global threats. Still, a big part of our potential to respond to global threats comes from socio-political decisions. Even in this field, open-mindedness, building trust, and showing solidarity are keys for successful future cooperation.

VLASTA ŠIKIMIĆ
University of Tübingen

Mathematical Philosophy

Deductive reasoning is widely viewed as the best and most characteristic basic source of justified mathematical belief. In typical cases, we take ourselves to have good grounds for accepting a mathematical claim P only if we (or someone) has a correct proof of P .

So it’s interesting that many unproved conjectures

are confidently believed to be true by large numbers of mathematicians— $P \neq NP$, the Riemann Hypothesis and Goldbach’s conjecture, to name a few. To the extent that these beliefs are justified, the relevant grounds are inductive considerations of one sort or another. The use of such methods raises some big epistemological questions: What kinds of inductive evidence are admissible in mathematics? How much weight do these kinds of evidence carry? Can mathematical facts ever be known on purely inductive grounds?



One of the simplest types of non-deductive reasoning is enumerative induction: the inference from “ $P(a_1)$ is true”, “ $P(a_2)$ is true”, “ $P(a_3)$ is true”, etc. to “ $\forall x P x$ is true” (where the a_i make up less than the whole domain). We have lots of enumerative evidence for some unproved conjectures. Goldbach’s, for instance—the statement that every even integer greater than 2 is the sum of two primes—has been checked for all cases up to $4 \cdot 10^{18}$, with no counterexamples found. How confident should this make us that the conjecture is true?

Our epistemic situation here is a bit strange. In one sense, this quantity of enumerative evidence is huge—we haven’t carefully checked 10^{18} instances of any empirical regularity, although we’re pretty confident that some of those are genuine laws. In another sense, though, 10^{18} confirmed instances is almost nothing. There are infinitely many even integers, so any finite number of verifications is the tiniest of drops in the bucket. Where do these powerful but opposed considerations leave us?

These issues have been discussed at least since Frege, who took a skeptical view of inductive evidence in mathematics. Part of his reasoning was that, in the natural world, any given point in space and time is as good as any other (provided the local conditions are relevantly similar). So if, say, Newton’s laws hold in our solar system, we can fairly safely infer that they hold everywhere. By contrast, “[p]osition in the number series is not a matter of indifference like position in space” (*Foundations of Arithmetic*, §10). Frege’s skepticism about enumerative induction seems to have extended to all forms of non-deductive mathematical reasoning.

Subsequent authors have taken a more nuanced and often more sanguine view. George Pólya’s classic *Mathematics and Plausible Reasoning* (1954) chronicles instructive examples from various parts of math (concluding, contra Frege, that “the role of inductive evidence in mathematical investigation is similar to its role in physical research” (vol. 1, viii)). James Franklin’s “Non-deductive Logic in Mathematics” (1987, *BJPS* 38, 1-18) is a Pólya-inspired philosophical defense of a similar thesis.

Zooming in again on enumerative induction in particular, we find ourselves approaching What’s Hot territory. The recent debate here begins with Alan Baker, “Is There a Problem of Induction for Mathematics?” (2007, in Leng, Paseau and Potter (eds.), *Mathematical Knowledge*, OUP). Baker takes a rather pessimistic view. For him, what’s troublesome about enumerative evidence on scales like $n \leq 10^{18}$ is that such numbers are *minute*, meaning they fall within the range that ordinary humans can represent using standard mathematical notation.

While the notion of a “small number” is vague, Baker takes minuteness to be relatively sharp. He also claims that minuteness makes a difference in some real mathematical cases. Unfortunately his lone example no longer holds up: Baker had claimed that an upper bound on the size of a certain number related to the Riemann Hypothesis is minute if RH is true but non-minute if RH is false, but a 2010 result of Saouter and Demichel showed unconditionally that this number is at most a little over 10^{316} . So it’s unclear that there really are interesting systematic differences between minute and non-minute numbers. Hence it’s unclear that we should be especially mistrustful of the enumerative evidence for claims like Goldbach’s conjecture.

In his “[Empiricism, Probability and Knowledge of Arithmetic](#)” (2014, *Journal of Applied Logic* 12, 319-348), Sean Walsh raises another sort of objection to Baker. Just as all the numbers checked in our Goldbach verifications have been minute, all of humanity’s mature scientific observations have been made within, say, the last ten thousand years. This can also be viewed as a serious sort of “smallness bias”—while ten thousand years isn’t infinitesimal compared to the age of the universe, it is extremely tiny. If we can’t rationally make generalizations on the basis of samples suffering from smallness bias, then our empirical generalizations would seem to be irrational too. But of course Baker wants to avoid that conclusion. So smallness alone apparently isn’t the problem.

The latest addition to this literature is Alex Paseau’s “[Arithmetic, Enumerative Induction and Size Bias](#)” (forthcoming, *Synthese*). Paseau’s main contribution is a taxonomy of types of size-based skepticism. If you think that samples consisting of smaller numbers always provide weak inductive evidence compared to samples consisting of bigger numbers, you’re a *c-skeptic*. If you think there’s some natural division between small and large numbers such that all-small samples provide inherently weak inductive evidence, you’re an *s-skeptic*. If you think the problem with all-small samples is that they’re unrepresentative of the natural numbers in general, you’re a *u-skeptic*. Paseau constructs a simple Bayesian-like model of mathematical credence and uses it to argue that *u-skepticism* is the most plausible of the three. This seems like the right conclusion. But it doesn’t tell us how large a spread is required for a representative sample, or how strong our enumerative evidence is in real cases of interest.

For more on these themes, readers attending the 2022 Eastern APA meeting may want to drop by the Philosophy of Mathematics Association group session featuring Alan Baker, Silvia De Toffoli and me. Our topic is non-deductive evidence in math, and Alan will discuss his latest work on enumerative induction, responding to Walsh and Paseau. So hot it it hasn’t even happened yet!

WILLIAM D’ALESSANDRO
MCMP, Munich

EVENTS

JULY

[EoA](#): Speaker Series on the Ethics of Argumentation, virtual, 2 July.

[D-M&ER](#): Difference-Making and Explanatory Relevance, online, 12–16 July.

[LoSy](#): Panhellenic Logic Symposium, Volos, Greece, 14–18 July.

[AiCI](#): Advances in Causal Inference, online, 30 July.

SEPTEMBER

[PROGIC](#): Combining Probability and Logic, Munich, Germany, 1–3 September.

[CSPS](#): Congress of the Society for the Philosophy of Science, University of Mons, Belgium, 8–10 September.

[CIAML](#): Causal Inference and Machine Learning, online, 10–11 September.

[VoAS](#): The Varieties of Anti-Skepticism, Pamplona, Spain, 15–17 September.

[SPoI](#): Science and Philosophy of Imagination Conference, University of Bristol, 16–17 September.

[ECSQARU](#): European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, Prague, 21–24 September.

OCTOBER

[SaR](#): Science and Responsibility: On the Role of Values in Science, Warsaw, 8 October.

[IJCLR](#): International Joint Conference on Learning & Reasoning, virtual, 25–27 October.

COURSES AND PROGRAMMES

Courses

[CE](#): Computability in Europe 2021: Connecting with Computability Tutorials, 5–9 July.

Programmes

[MA IN REASONING, ANALYSIS AND MODELLING](#): University of Milan, Italy.

[APHIL](#): MA/PhD in Analytic Philosophy, University of Barcelona.

[MASTER PROGRAMME](#): MA in Pure and Applied Logic, University of Barcelona.

[DOCTORAL PROGRAMME IN PHILOSOPHY](#): Language, Mind and Practice, Department of Philosophy, University of Zurich, Switzerland.

[DOCTORAL PROGRAMME IN PHILOSOPHY](#): Department of Philosophy, University of Milan, Italy.

[LOGICS](#): Joint doctoral program on Logical Methods in Computer Science, TU Wien, TU Graz, and JKU Linz, Austria.

[HPSM](#): MA in the History and Philosophy of Science and Medicine, Durham University.

[MASTER PROGRAMME](#): in Statistics, University College Dublin.

[LoPhiSC](#): Master in Logic, Philosophy of Science and Epistemology, Pantheon-Sorbonne University (Paris 1) and Paris-Sorbonne University (Paris 4).

[MASTER PROGRAMME](#): in Artificial Intelligence, Radboud University Nijmegen, the Netherlands.

[MASTER PROGRAMME](#): Philosophy and Economics, Institute of Philosophy, University of Bayreuth.

[MA IN COGNITIVE SCIENCE](#): School of Politics, International Studies and Philosophy, Queen’s University Belfast.

MA IN LOGIC AND THE PHILOSOPHY OF MATHEMATICS: Department of Philosophy, University of Bristol.

MA PROGRAMMES: in Philosophy of Science, University of Leeds.

MA IN LOGIC AND PHILOSOPHY OF SCIENCE: Faculty of Philosophy, Philosophy of Science and Study of Religion, LMU Munich.

MA IN LOGIC AND THEORY OF SCIENCE: Department of Logic of the Eotvos Lorand University, Budapest, Hungary.

MA IN METAPHYSICS, LANGUAGE, AND MIND: Department of Philosophy, University of Liverpool.

MA IN MIND, BRAIN AND LEARNING: Westminster Institute of Education, Oxford Brookes University.

MA IN PHILOSOPHY: by research, Tilburg University.

MA IN PHILOSOPHY, SCIENCE AND SOCIETY: TiLPS, Tilburg University.

MA IN PHILOSOPHY OF BIOLOGICAL AND COGNITIVE SCIENCES: Department of Philosophy, University of Bristol.

MA IN RHETORIC: School of Journalism, Media and Communication, University of Central Lancashire.

MA PROGRAMMES: in Philosophy of Language and Linguistics, and Philosophy of Mind and Psychology, University of Birmingham.

MRES IN METHODS AND PRACTICES OF PHILOSOPHICAL RESEARCH: Northern Institute of Philosophy, University of Aberdeen.

MSc IN APPLIED STATISTICS: Department of Economics, Mathematics and Statistics, Birkbeck, University of London.

MSc IN APPLIED STATISTICS AND DATAMINING: School of Mathematics and Statistics, University of St Andrews.

MSc IN ARTIFICIAL INTELLIGENCE: Faculty of Engineering, University of Leeds.

MSc IN COGNITIVE & DECISION SCIENCES: Psychology, University College London.

MSc IN COGNITIVE SYSTEMS: Language, Learning, and Reasoning, University of Potsdam.

MSc IN COGNITIVE SCIENCE: University of Osnabrück, Germany.

MSc IN COGNITIVE PSYCHOLOGY/NEUROPSYCHOLOGY: School of Psychology, University of Kent.

MSc IN LOGIC: Institute for Logic, Language and Computation, University of Amsterdam.

MSc IN MIND, LANGUAGE & EMBODIED COGNITION: School of Philosophy, Psychology and Language Sciences, University of Edinburgh.

MSc IN PHILOSOPHY OF SCIENCE, TECHNOLOGY AND SOCIETY: University of Twente, The Netherlands.

MRES IN COGNITIVE SCIENCE AND HUMANITIES: LANGUAGE, COMMUNICATION AND ORGANIZATION: Institute for Logic, Cognition, Language, and Information, University of the Basque Country (Donostia San Sebastián).


OPEN MIND: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.

RESEARCH MASTER IN PHILOSOPHY AND ECONOMICS: Erasmus University Rotterdam, The Netherlands.

LogiCS: Joint doctoral program on Logical Methods in Computer Science, TU Wien, TU Graz, and JKU Linz, Austria.

Jobs

ASSISTANT PROFESSOR IN: Applied Probability, University of Nottingham, deadline 14 July.



JOBS AND STUDENTSHIPS

Studentships

DOCTORAL PROGRAMME IN PHILOSOPHY: Language, Mind and Practice, Department of Philosophy, University of Zurich, Switzerland.