Eesti Haridusteaduste Ajakiri, nr 9(1), 2021, 33–59
https://doi.org/10.12697/eha.2021.9.1.02b

Teaching mathematics as a way of thinking – not calculating

Keith Devlin[a1]

[a] *Stanford University (Emeritus)*

At the end of 2018, I retired from my regular university position after a fifty-year career as a mathematician. During that entire career, I used almost ***none*** of the mathematical techniques I learned at school or in my undergraduate mathematics degree. And as I show below, my experience is by no mean atypical for people whose careers involve regular use of mathematics.

That's not to say my time learning mathematics at school and university was wasted. Quite the opposite. The point is, *what* I learned (and practiced and mastered over thousands of hours) was of little direct consequence. What mattered was *the way of thinking* I acquired during my school and undergraduate education. The valuable skill I took away from my school and bachelor's education was not the ability to execute a range of mathematical procedures, but mastery of a particular way of thinking: what some of us have referred to as *mathematical thinking*.

Important aspects of mathematical thinking are exploring, questioning, working systematically, visualizing, conjecturing, explaining, generalizing, justifying, and proving (but excluding the execution of formal procedures either done by machines or viewed as a "lower-level", mechanical activity). See, for example, Stacey (2006); Devlin (2012a,b,c); Singh et al. (2018); NRICH (2020).

Mathematical thinking is what this essay is about. But before I start, it should be noted that I write from the perspective of a career that spanned both academic research in pure mathematics and the world of applied mathematics, where I worked on a wide range of real-life problems for private industry and government.

[1] Graduate School of Education, Stanford University, 450 Serra Mall, Stanford, CA 94305, USA; kdevlin@stanford.edu

Mathematics as a tool – the technological revolution

In the second half of my career, I used mathematics as a tool to solve real-world problems (and at the same time helped develop new mathematical tools). In doing so, I experienced a dramatic shift in the way mathematics is done (i.e., used).

Most of that shift took place in a twenty-year period from 1970 to 1990. Though many people, including parents and teachers, are even to this day, unaware there has been a change, since 1990, practically everyone who uses mathematics professionally does it *in a totally different way* than previously. It's all due to technology, and it is no exaggeration to say it was a *revolution* in mathematical praxis.

The first step towards this revolution took place in the early 1960s, with the arrival of the electronic calculator, making it no longer necessary for professionals to carry out numerical calculations by hand. The electronic calculator – the early ones were expensive, desktop devices – could perform calculations far quicker than humans, and without errors.

The first electronic calculators were expensive (around $5,000 each). They were used almost exclusively in the workplace, alongside (and eventually replacing) heavy, electrically-driven mechanical calculating machines that had been available since the early twentieth century – and which in turn had replaced still earlier, hand-powered desktop calculators (See Fig. 1).

Figure 1. Calculators going back in time. (TL) The still-popular Texas Instruments TI84 hand-held graphing calculator, introduced in 2004 (cost around $100); (TR) Hewlett Packard's HP9100 electronic desktop calculator, introduced in 1965 (cost $5,000); (BL) the Marchant electrically driven mechanical desktop calculator, ca.1920s; (BR) hand-operated desktop calculators from the late nineteenth century. Images from Wikipedia.

But then, at the start of the 1970s, small, cheap, pocket-sized electronic calculators became available, making the new technology available to all (at least in the developed world). In the workplace, the electronic pocket calculator rapidly ruled supreme. By 1973, the desktop calculator business had died. People who worked with numbers simply stopped doing hand calculation. Nevertheless, schools continued to teach arithmetic skills. As I will explain later, they did so for the wrong reasons, and in consequence, they taught it inappropriately, but their instinct that it was still important to learn arithmetic was entirely correct.

Because of the electronic calculator, when I arrived at university to study mathematics in 1965, I knew I would never again need the fluency at arithmetic I had developed through many years of school education. On the other hand, I did have to spend a great deal of my time as an undergraduate mathematics major mastering a whole range of algorithms and techniques for performing various kinds of numerical and symbolic calculations, geometric reasoning, algebraic reasoning, and equation solving. In order to solve mathematics problems, I had to be able to crank the algorithmic and procedural handles. There was no other way. There were no machines to do it for me the way the calculator could perform arithmetic calculations.

That remained the case after I graduated, and throughout the early part of my career as a mathematician. Then, in 1985, the arrival of widely available, mass-produced *graphing* calculators extended the reach of digital technology well beyond arithmetic, providing tools that would perform algebraic derivations and draw graphs. Because they were relatively cheap (under 100 Euros), graphing calculators led to a revolution in how school mathematics could be taught, particularly in science education classes. (With good reason, mathematics teachers continued to insist on some hand calculation.)

Soon after the graphing calculator revolutionized school science education, a far greater technology revolution swept over university mathematics education and changed forever the way professionals did the mathematics. In June 1988, mathematician Stephen Wolfram released the first version of his massive package *Mathematica* (https://www.wolfram.com/mathematica/). (See Fig. 2.)

Quite simply, *Mathematica* can execute *any* mathematical procedure, in *any* branch of mathematics.

Soon after Wolfram launched his product, Canadian developers released a similar system called *Maple* (https://www.maplesoft.com), and a number of other analogous products came out. These products did for almost all procedural mathematics what the electronic calculator did for arithmetic: they made the mastery of procedures obsolete as a human skill. (*Obsolete except for educational purposes*, I should add, of which more later.)

For the first time in history, being able to perform calculations, or execute *any* mathematical procedure, fluently and accurately was no longer a necessary requirement for using mathematics. This highlighted the distinction, which was always there but had been invisible to most non-mathematicians, between the *routine* parts of using mathematics (executing procedures) and the *creative* parts.

Corporate Speakers at the *MATHEMATICA*™ Product Announcement (June 23, 1988):

- **Forest Baskett**, Vice President, Research & Development, Silicon Graphics Computer Systems
- **Gordon Bell**, Vice President, Research, Development and Engineering, Ardent Computer
- **Steven Jobs**, President NeXT, Inc.
- **William Joy**, Vice Presiders, Research & Development, Sun Microsystems
- **Vicky Markstein**, Research Staff Member, IBM
- **Eric Lyons**, Director of Technology, Autodesk, Inc.
- **Larry Tesler**, Vice President, Advanced Technology, Apple Computer
- **Stephen Wolfram**, President, Wolfram Research, Inc., and Professor of Physics, Mathematics and Computer Science, University of Illinois

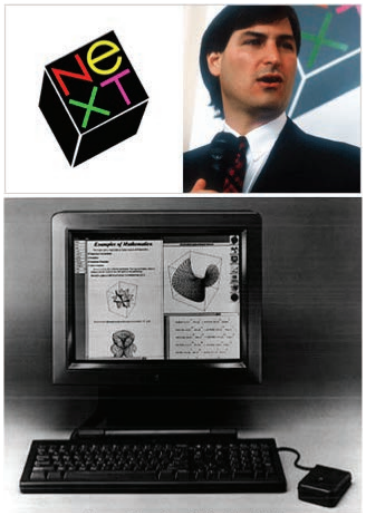

Bundled with NeXT

Figure 2. In June 1988, Professor Stephen Wolfram introduced his new computer algebra system *Mathematica*.

Orchestras of learning

For a few years, products like *Mathematica* and *Maple* were used mainly in university departments of mathematics, physics, and engineering. They were expensive, challenging to use, and ran only on upper-end personal computers. As a result, few schools made use of them, and the new systems had minimal impact on school mathematics and science education. (I know. I was a member of Wolfram's Scientific Advisory Board in the early days of the company.)

That changed dramatically in 2009, with the release of *Wolfram Alpha*, which made the power of *Mathematica* available in a Cloud-based application that could be accessed (for free) on the Web from any PC, tablet, or mobile phone. Moreover, *Wolfram Alpha* had a simple user interface that made it possible to execute any mathematical procedure with as much ease as using an electronic calculator.

The easiest way to get a sense of how *Alpha* works is simply accessing it on the Web and explore it for a while (http://wolframalpha.com). The point I want to make is that it made it possible for people to use mathematics without having expertise in executing any particular procedure. (I'll come later to exactly what knowledge is required to do this. It is considerable.)

The arrival of *Wolfram Alpha* has changed forever *the way* people can use mathematics. Today, mastering procedures is no longer the price everyone has to pay to use mathematics.

Two years after *Wolfram Alpha* came out, in 2011, Eli Luberoff released the Web-accessible and mobile-based system *Desmos*, a very powerful graphing calculator designed for school mathematics education. It could, and in many schools did, totally revolutionize the way school mathematics is taught and done, in the same way *Alpha* revolutionized how professionals use mathematics. (See, for example, Abramovich 2013.)

In both cases, the new tools shifted the emphasis in mathematics from executing procedures – which had dominated mathematics learning and used throughout its entire five-thousand-year history – to creative thinking and problem-solving. This was a major shift. Yet it happened so rapidly, most people outside the worlds of science, engineering, and mathematics had no idea it had occurred. In the United States, this was illustrated dramatically with the adoption, in 2010, of the Common Core State Standards for Mathematics Education (http://www.corestandards.org/Math), which laid out what school students throughout the nation should know at the conclusion of each school grade.

Aimed at ensuring that American students would graduate from school equipped to use mathematics effectively in today's world, both as citizens and in their work (whatever that may turn out to be), the Common Core standards met with significant opposition, not all of which has evaporated in the years since. The problem was that the committee of experts who drew up the standards was fully aware of the way technology had revolutionized how mathematics is done today and what that meant for the way mathematics should be taught. Still no one seemed to fully realize the degree to which society as a whole, and many parents (and some teachers) in particular, were totally unaware of that revolutionary shift.

To help people understand what it is like to use mathematics in today's world, I often draw an analogy with the world of music. Being a mathematician before 1990 was like mastering many instruments in an orchestra: the arithmetic instrument, the geometry instrument, the trigonometry instrument, the algebra instrument, the calculus instrument, and so on. The more mathematical instruments you mastered, the greater your power as a mathematician.

In contrast, using mathematics today is more like being a conductor of the orchestra. To conduct that orchestra well, you have to know what all of the instruments are capable of, and you surely need to gain some experience with a few of them, at least one of them fairly well. But there is no need to be world-class in any of them. The individual musicians in the orchestra "do all the detailed work". As a conductor, you have to know how and when to make them work together, directing which instrument(s) to use for each purpose as you progress through the symphony.

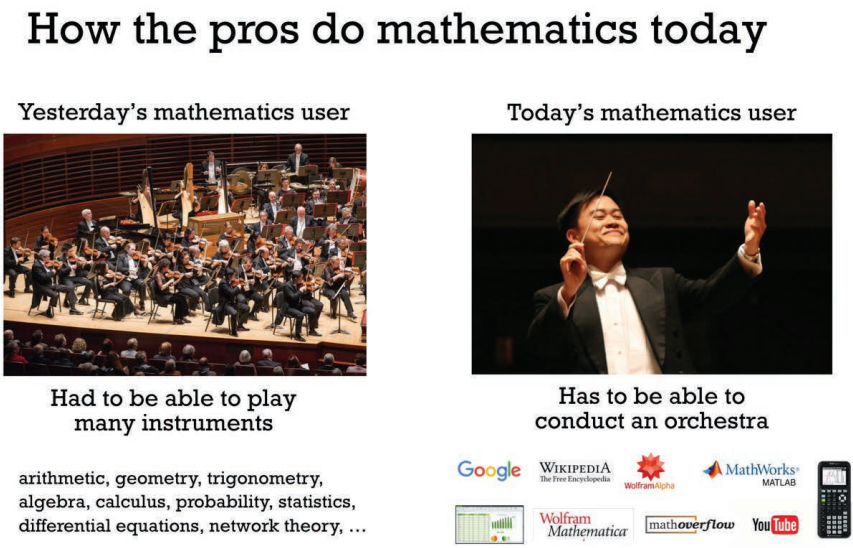

Figure 3. Orchestras of learning. See text for discussion. Public domain images.

Clearly, with mathematics being done that way, the experience of using mathematics is very different than it was throughout the entire previous history of mathematics. In particular, gone is the need to be good at any kind of calculation. Mathematicians today do not need to be able to calculate quickly or accurately; indeed, *they almost never do that*. The detailed execution of any formal procedure or algorithm is now done by machines. The machines do it considerably faster than humans ever could, they make far fewer errors (essentially none), and they do it with far bigger data sets than the human brain could handle.

The rise of mathematical thinking

One curious aspect of the opposition of American parents to the Common Core was highlighted by posts on social media by scientists, engineers, and other mathematical professionals, criticizing the way their children were being taught at school. Though those parents knew fully well how today's professionals use technology to execute mathematical procedures, they felt it was important for their children to spend the school years mastering the traditional hand methods – the so-called "standard algorithms." Though experience doing hand calculation is a valuable, indeed essential, step towards proficiency in mathematical thinking, those parents' knowledge of mathematics education (a subject grounded in decades of classroom research) was at best minimal.

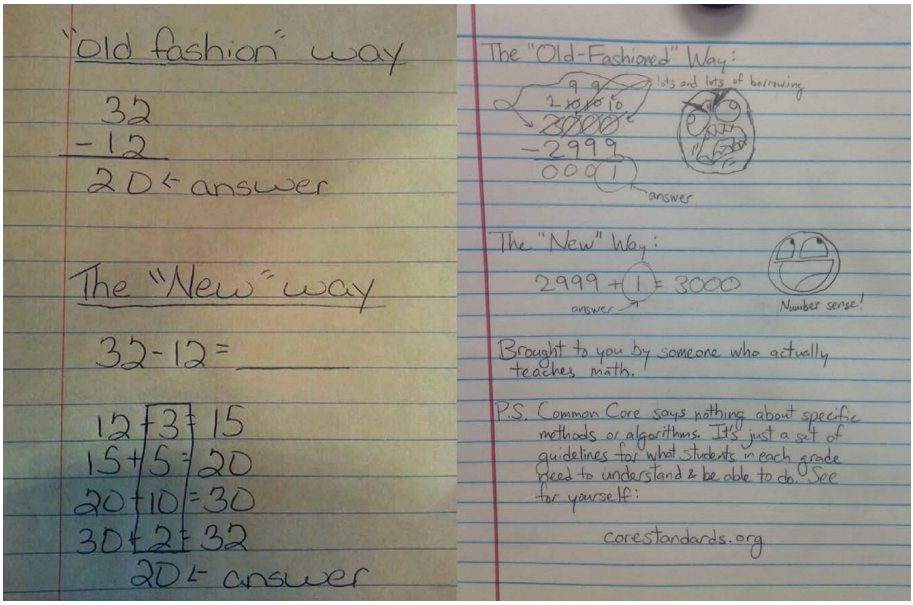

Figure 4. The image on the left was posted on Facebook by a parent of a student being taught in alignment with the Common Core. A teacher's response is shown on the right. See text for discussion. Social media images.

An illustration of one such parent critique is shown in Fig. 4. The parent erroneously believed the child was being taught an inefficient way to calculate. As a teacher responded on the right, the child's work did not display an algorithm being followed. Rather, the child was *reasoning creatively with* the place-value number system to explore (and hence understand) how subtraction operates. The educational goal was not the computation of a number – a calculator can do any computation in an instant – but the development of *number sense*, a hugely important ability in the age of ubiquitous computing resources.

Briefly, number sense is *fluidity and flexibility with numbers, the sense of what numbers mean, and an ability to perform mental mathematics and to look at the world and make comparisons.* For example, if you were faced with having to calculate 16 × 24, how would you proceed? Someone with number sense might reason like this: It is easy to calculate 16 × 25 since that's the same as 4 × 100, which is 400, but that is one too many 16's, so I need to subtract 16, which gives me 384. Being able to reason about numbers like that is an example of number sense.

The point is, most of today's professionals learned their mathematics the traditional way because there was no alternative at the time. But there are four problems with that historical pathway:

1) Wasted time. It means spending a great deal of time learning and practicing methods that were honed over centuries to be computationally efficient for hand calculation, with a historically-necessary emphasis on speed and accuracy, neither of which is important when, in any real-life situation, a machine will be used.
2) Lack of student understanding. By emphasizing computational efficiency, those methods do little to help students *understand* the procedure in *mathematical* terms. Indeed, steps often focus on how to lay out the calculation on a page, which can be at most tangentially related to the underlying structure of the number system and its arithmetic. Mastery of procedures without understanding is brittle knowledge that can be effectively applied only in narrowly constrained circumstances (like exams).
3) Lack of understanding later in life. Lack of student understanding can prove an obstacle when, later on in life, those who become scientists and engineers are faced with having to engage in serious mathematical thinking to make creative and effective use of the powerful new technologies they use. Indeed, this was demonstrated repeatedly by their very social media posts, which showed that those parents were totally unable to understand the conceptually based methods for calculation that their children were being taught. The fact that those professionals *could not follow* a (novel, to them) procedure for performing a calculation in elementary arithmetic shows that, although they could *mechanically* carry out computational steps they had learned, they had never achieved even a basic level of *understanding*.
4) Negative attitudes. While parents who were scientists or engineers had most likely enjoyed their school mathematics classes, and had taken pride in mastering the methods (as I did), almost certainly the vast majority of their fellow students became frustrated and developed a negative attitude to mathematics, often giving it up at the first opportunity. This is a significant loss both to them, in terms of future career opportunities, and to society in terms of wasted potential talent.

So those Common Core critics were wrong. But they were right in one respect. Learning to reason with numbers and other mathematics abstractions is a valuable, indeed essential part of mastering mathematical thinking. Where they went wrong was in tacitly assuming a goal that was out-of-date. In today's world, there is no need to teach calculation, and more generally, any procedural mathematics, *for execution*. Machines do that far better than we can. Rather *the goal is to achieve sufficient understanding* to be able to make good use of those machines in using mathematics to solve (real-world) problems. In terms of my music analogy, we need to teach the next generation to be conductors of orchestras, not masters of individual instruments.

Sure, to become a good conductor, you need to achieve a sufficient mastery of at least one instrument, maybe more. But no conductor masters them all, nor do they need to. The skillset for conducting is surely grounded in learning to play one or more instruments, but it is a very different skillset than that required to play, say, a violin or a piano. The same is true for mathematics.

What mathematics should we teach?

As with music, it does not really matter what mathematical instruments (areas of mathematics) you master. It may as well be arithmetic, algebra, and geometry since they are all entry-level subjects, and all are relevant in many careers and various walks of life. Moreover, all of today's teachers are familiar with those subjects.

To that list, we should surely add data science (which includes the topic of algorithms) since that subject plays such a major role in today's world. To give just one, highly topical example, the effective use of representations as in Fig. 5 requires the ability to correctly interpret the relative shapes of the different scenario-curves, understanding what they each signify, in order to make better decisions. In the early days of the pandemic, US decision-makers did a poor job of reading the data, and the result was a death count an order of magnitude larger than it should have been. Good data science skills are crucial in modern society.

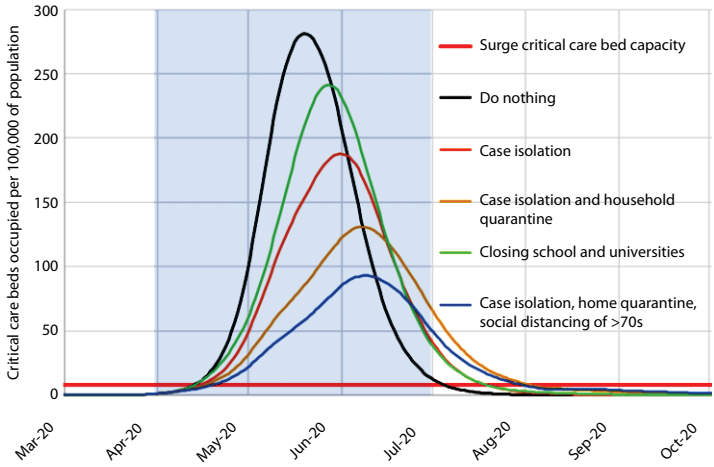

Figure 2: Mitigation strategy scenarios for GB showing critical care (ICU) bed requirements. The black line shows the unmitigated epidemic. The green line shows a mitigation strategy incorporating closure of schools and universities; orange line shows case isolation; yellow line shows case isolation and household quarantine; and the blue line shows case isolation, home quarantine and social distancing of those aged over 70. The blue shading shows the 3-month period in which these interventions are assumed to remain in place.

Figure 5. The importance of data science. This figure is taken from a government planning document used by the UK government in early March to explore various mitigation strategies for the coronavirus epidemic that was just beginning. The actual figures are estimates produced by a mathematical model based on existing techniques of mathematical epidemiology. (Ferguson et al. 2010)

For algebra, linear equations and inequalities in two or more unknowns are hugely important in today's world, which provides an argument for building an algebra course around that topic.

The main use of that algebra is in formulating and solving optimization problems. Google, shipping companies like UPS and FedEx, major airlines, and large online retailers all make heavy use of optimization using linear algebra.

To be sure, the optimization problems those companies depend on typically involve thousands or even millions of unknowns, which is well beyond human capabilities. Computer packages are used to solve them. No human could ever cope with that. But for linear optimization, working on examples with just two or three unknowns provides a good understanding of the method. (This is not true for all mathematical topics, but it is here.)

So, given the frequency with which solutions of linear equations and inequalities crop up in a great many of today's real-world problems that affect our lives, there is a good case to make for teaching methods to solve, by hand, one and two variables, examples.

In contrast to the above subjects, calculus is most definitely not necessary at the school level. For one thing, it is only important for students who wish to pursue science, engineering, or mathematics itself. (Data science is important for everyone.) But also, calculus cannot be taught well at the high school level since it is considerably more sophisticated than anything else in the school curriculum. Students who do take calculus at school frequently have trouble with university calculus courses later on, since their school experience leaves them with a superficial, and essentially procedural, understanding that gives them a false sense of security during the first weeks, which eventually gives way to an unpleasant, and occasionally disastrous, experience when they find their weak understanding is inadequate for the more advanced part of the course.

As I'll stress again, the point is that while *what* we teach is subject to some debate, what should not be up for discussion is *why* we teach a particular topic, and *how* we teach it. The answer to the second question ("how") depends on the answer to the first ("why").

And the reason there should be no debate is that successful mathematics education should ensure that future generations are able to make effective use of mathematics in the world they will inhabit, and in that world, *mathematical thinking* is the key skillset. (Computers *calculate* and *execute procedures*; people *think*.)

By the way, number sense is a part of mathematical thinking, so too is the capability to reason logically from assumptions to conclusions.

Remember, it is not mathematics that has changed in the digital age, though there have been changes in the form of new branches of mathematics that resulted from the growth of computer technology. (Fractal geometry, for example.) What has changed is the way people use mathematics. With the change in praxis from the calculation and the execution of procedures to mathematical thinking, has come to a change – or rather, there is an emerging process of change – in what is required of an education system to produce effective users of mathematics (a.k.a. *mathematical thinkers*).

Being able to calculate quickly, efficiently, and accurately used to be essential. Now, it is not required. In place of that skillset (which took most people considerable time and effort to master, with many dropping by the wayside in the process) is a new set of skills. Those new skills – mathematical thinking – are, in fact, much closer to those in the humanities or the creative arts than most people yet realize or, in some cases, are willing to contemplate.

To be sure, to enjoy mathematics, you have to be intrigued by the very idea of formally specified abstractions and context-free, formal reasoning. Not everyone will see mathematics as having appeal, even if they can do it. But then, few among us can see the attraction in everything our fellow humans

decide to pursue. From the human perspective, it's not so much that today's digital mathematical tools have added something to the discipline; rather, they have removed what for many was an obstacle.

To summarize, since the goal of school and early college-level mathematics education today is mathematical thinking ("conducting the orchestra"), exactly ***what*** mathematics we teach is secondary to ***how*** we teach it. Given today's technologies, we are certainly not constrained, as we were in the past, to teaching mathematics that can be done when students have to perform all the calculations by hand. On the other hand, there is certainly an educational benefit from performing calculations by hand, since that is how we come to understand numbers and the *process* of calculation. Likewise, there is benefit from becoming familiar with algebraic abstractions and having some experience with solving simple linear equations.

How do we teach mathematical thinking?

The suggestions for mathematical topics to teach I gave above are just that: suggestions. As I said earlier, the issue is *how* we teach, not *what* we teach. At its heart, mathematical thinking is the same, whatever the topic. To re-use my music analogy, all the members of the orchestra play *music*, even though the individual instruments look and sound different. Music is music.

It is possible to develop proficiency in mathematical thinking by focusing just on arithmetic (whole numbers and fractions). This point was made effectively by the Chinese-born mathematics educator Liping Ma (2013). Ma based her case on teaching done in China, where she began her teaching career before moving to the United States and obtaining a doctorate at Stanford University.

A method of teaching that works in China may not work as well in more open Western societies. In any case, there is no need to be so restrictive. Teaching arithmetic, algebra, and geometry, together with data science, provides a broader perspective and facilitates a greater variety of examples, all self-evidently highly relevant to everyday life, which can help motivate students. With the range of technological aids available today, students can work on real-world problems of relevance to their lives, free from the constraint that governed mathematics education in past centuries that classroom problems had to be solvable using hand methods that were within the capabilities of the students.

As to the "how" do we teach, this is where society will sooner or later have to face the reality that developing mathematical thinking ability requires a very different conception of teaching than the one most familiar to teachers. Except that teachers *are* familiar with it, just not in the mathematics classroom.

In the early stages of learning, most school disciplines involve acquiring a substantial amount of information. To progress in the subject, the student first has to *know* quite a lot. It involves a lot of teacher instruction, reading, and these days, (but not when I was a student) video watching. Tests and exams typically ask the student to demonstrate that they have assimilated enough of that information.

Mathematics is an exception. Yes, there are some facts to be learned, but mostly, learning mathematics involves *solving math problems*. Look at any mathematics textbook. After a short passage presenting some information, the student is presented with a much longer section that provides examples of the thinking required to solve problems using that new information, followed by a long list of problems for the student to attempt. Compare that with a textbook on introductory biology or physics, say, or history, social studies, geography, or literature.

Learning mathematics is primarily about *doing,* not about *knowing,* and is more akin to sports, music, or handicrafts (or the laboratory sections of science classes, which typically follow and depend upon fact-learning classes). What role does the teacher play in these disciplines?

As a teacher or a parent (the target readers of this essay), think back to when you took lessons in, say, driving, playing a musical instrument, tennis, golf, skiing, chess, creative writing, or speaking a foreign language. What role did the instructor play?

"Coach" is surely the best way to describe it. You learned how to perform that activity by doing it under the watchful eye of an expert who was able to guide you towards improvement. That's how we best learn how to *do* something. Yes, we may be able to make progress on our own by attending lectures, reading books, or watching videos. But it generally goes much faster, and we achieve far better levels of performance when our learning is guided by an expert coach.

As a *doing* subject, mathematics is also best learned that way. One of the reasons for using the term *mathematical thinking* is to emphasize that it is primarily a *doing* subject. The doing involved is active, and frequently creative, *thinking*.

At present, the coaching-style approach to mathematics learning is standard only at the university doctoral education level. A few universities also provide it systemically at the undergraduate level; for example, Oxford and Cambridge in the UK, with their "Tutor system".

Knowing the value of the coaching approach, many college and university faculties try to find time to provide coaching sessions, as do some school mathematics teachers. But with systemic school education structured the way it

is, this is hard to do. When a teacher focuses on one particular student, as good coaching requires, they cannot be attending to the needs of the others in the class. Splitting a class into small groups can go some way towards the desired goal, as can the use of teaching assistants. But that is not ideal.

And this brings me to one of the main points I want to make. Throughout most of the world the school systems adopted a "production-line" approach, modeled on and designed to prepare future citizens for the industrial age of the nineteenth century, when they were developed and deployed. In the case of mathematics, curricula and textbooks were designed to support that model. The main societal need was for an arithmetically able workforce to support, in particular, the mechanized, production-line manufacturing that drove the Industrial Revolution. There was little need for individual creativity. Fast, efficient, accurate rule-following, with everyone doing things the same way, was the order of the day.

Today, things are very different. Anything that can be done by routine rule-following is now done by machines. The primary need today is for creative thinkers and problem solvers to do the things the machines cannot. And the reality is that the very automation that made redundant the human skills of the nineteenth century has given rise to a data-rich society where there is a great demand for such (human) skills. Most people involved in education know this, of course. What makes it difficult to change education to meet the need is that it requires a major restructuring of the entire education system, including the structure of the teaching profession.

My own view is that school mathematics classes need (at least) two instructors. One will be the classroom teacher, who is in charge of the entire class, and the other (and maybe it requires more than one) would be a "tutor" who, on a rotating schedule, spends some time each week with each student on a one-on-one basis. (Maybe twice a week, with additional sessions on an as-needed basis if the teacher thinks it necessary.) The tutor could very likely operate remotely, over a video link with a shared workspace.

Ideally, the teacher would have a bachelor's degree in mathematics and an additional credential in mathematics pedagogy. The tutor must be someone with a broad knowledge of mathematics who has used mathematics extensively in their career, possibly a retired scientist or engineer. While the tutors would certainly require some training in pedagogy and how people learn mathematics, the primary skill they would bring is a deep and broad knowledge of mathematics, as it is practiced in the world today. Few teachers have such experience—teaching is a demanding career on its own.

In major conurbations, there is surely a good supply of such individuals. In more sparsely populated regions, video-links would be the way to go. Indeed,

the tutor's role might best be served by an "Uber-like" gig-economy system that connects classroom teachers to tutors with particular areas of expertise.

In any event, I put this suggestion out there for reflection and discussion. I should stress that I write from the perspective of a professional mathematician whose teaching experience is mostly, though not exclusively, gained in colleges and universities, with some brief excursions into high schools and occasionally middle schools. Thus, my suggestion is driven by what might best serve society (which, of course, is made up of the people in it) and a strong belief that there are really not many options open to us if we are to provide future generations with the mathematics education they, and hence society, need. Though I have been involved in education studies for several decades and have likely read considerably more about mathematics education research than many excellent teachers, I do not have extensive experience teaching at the school level. Those who do have such experience may be able to bring other perspectives to bear.

I am aware that middle-class parents of sufficient financial means frequently ensure that their children have regular coaching from a tutor by hiring freelance tutors to come to their homes once or twice a week and work with those children. As a graduate student, I supplemented my state maintenance grant with income from private tutoring. This kind of coaching should be available to all students.

The professional mathematician's toolbox

The proposal (for discussion) I made above of a systemic change in school mathematics education that I believe the national need requires was based on my decades of long experience using mathematics in a variety of contemporary domains. As I have indicated already in this essay, hardly any of the work I did involved me solving (by hand) an equation, or indeed anything that you might observe students doing in the traditional math classroom. Rather, as the technologies available grew over many years, I evolved a way of working on new problems.

I'll describe that approach in a moment. First, I should make it clear that I am talking about real-world problems here, not the stylistic "mathematics problems" you find in traditional mathematics textbooks, which are designed to provide exercises in executing a particular procedure. I'll come back to this distinction in the next section.

The problems I worked on arose in a number of domains where I was hired as a contract researcher. I contributed to projects for a large electronics company in the UK, for a European civil construction conglomerate, a US production-line manufacturer, the US intelligence service, and the US Navy and

Army. The lesson I learned time and again was that my value to the project did not lie in my ability to solve intricate mathematics problems. That was rarely required. Rather, what I contributed, and what I was repeatedly told was the benefit I brought, was the way I approached the problem. I provided mathematical thinking.

Indeed, those experiences are what led me to start using the term "mathematical thinking" to refer to that way of working, leaving people to continue to view "mathematics" as the stuff of the school mathematics class.

Fig. 6 captures the way I approach each new problem.

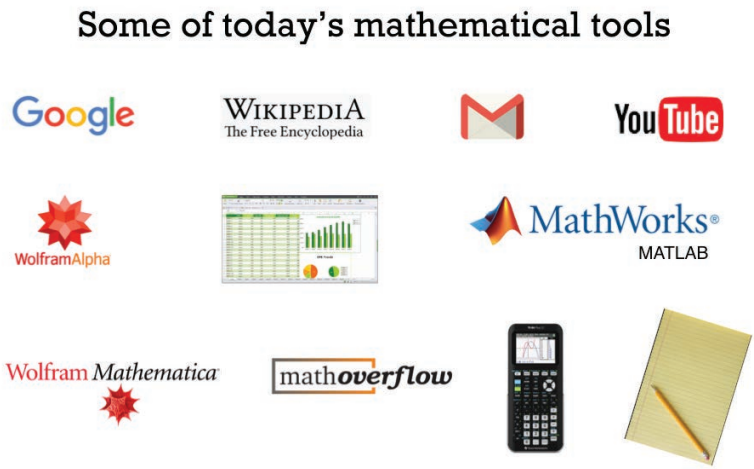

Figure 6. A typical mathematician's toolkit today. This is my own, with three tools added for pedagogic reasons. I normally use the tools in the order they read starting top-left.

I should stress that this is my own particular methodology. Other mathematicians and professional users of mathematics (such as physicists and engineers) who have seen this list have said they have a similar one, but not identical. We are each guided by the particular problems we have worked on. In fact, the figure includes some items that I rarely if at all use. For expository purposes, I expanded it to include tools others told me they use. Generally speaking, I use the tools in the order they appear in the figure, going left-to-right, line by line. Let me tell you why and how I use them.

Google: First up is *Google*. There are few problems that someone else has not already looked at. A quick *Google* search can rapidly provide you with links to relevant work done by others. This can prevent you from putting time and effort into duplicating work already done and can also provide hints as to how you should proceed.

The skill required here – and it is considerable – is knowing which key term(s) to search on. If a mathematician has been hired to work on a problem, it is surely one that the employer has been unable to solve on their own. So you can assume they think it requires mathematical expertise. It's possible you, as the mathematician, will end up telling them that mathematics is unable to help. I have certainly been in that position. Science fiction novels and movies, in particular, have led mathematical laypersons to approach me with some (currently) very unrealistic assumptions, particularly when artificial intelligence is concerned.

But, almost certainly it will be up to you to determine what kinds of mathematics will help and to figure out how to frame the problem, or part of the problem, to make that possible. The more information you can provide *Google*, the more likely it will be that it returns useful leads you can follow. Critics who say using tools like *Google* are "educational cheating", or that hiring a consultant is a waste of money if they simply use *Google* to find the solution to the problem (and I have heard both) are totally unfounded. The skill required here is knowing what to search for. That skill can come only through experience.

Wikipedia: My initial *Google* search sometimes returns links to papers that report research using mathematical techniques that I have never heard of or have encountered (and maybe even used) myself long ago but can no longer remember the details. *Wikipedia* usually fills in the gaps in my knowledge or memory. It is usually an excellent, reliable, and up-to-date source for advanced mathematical topics, since only experts can contribute.

Email: Since I have been in the mathematics business for many decades, I know, and know of, a wide range of experts in different areas of mathematics, and occasionally I email them when I have a specific question I am sure they will be able to answer – and if they don't, likely no one will. This typically works only for very specific, detailed questions, and is a useful resource only among a personal network built up over many years working in the field.

YouTube: When a source refers to a mathematical technique I may be able to use, but is new to me or long forgotten, I can usually find one or more instructional videos on YouTube that quickly get me up to speed. Again, skill is required in searching on the correct terms, and you need to be able to evaluate the reliability and accuracy of the video. These are general Web skills we all need in order to make safe and effective use of the Web as a resource.

Wolfram Alpha: Armed with a general sense of what mathematical techniques I should bring to bear, the next step is to make some initial attempts to formulate and explore the mathematical aspects of the problem I have identified. At this stage finding a solution is rarely the goal. This is still very much exploration to better understand the problem. *Alpha* can be particularly useful.

Spreadsheet: A spreadsheet program like Microsoft *Excel* is also useful in making an initial exploration of the problem. The power of packages like *Alpha* and *Excel* means that if I wanted I could plug in some real data, but at the early stages, I often use simplified data, or even very simple data I make up, as I seek to get an overall understanding of the problem.

MATLAB: I have never used MATLAB, but many other professionals do, so I include it here.

Mathematica: I have used *Mathematica*, and indeed was on Wolfram's Advisory Board in the early years of the program, but these days I find *Alpha* meets my needs. But that likely reflects the nature of the projects I am offered and accept. (Word gets around that I can be useful with certain kinds of problems.)

math*overflow* is a Web resource inspired by the computer engineers' stack*overflow*, where people can pose nerdy, technical questions online and nerdy, technical people around the world can give nerdy, technical responses. The level of sophistication of both questions and answers varies enormously, but it can be very effective. I have browsed it occasionally out of curiosity, but never used it in my work.

Graphing calculator: With *Alpha* to hand, I never make use of a graphing calculator, but millions do, and in many parts of the world this and a smartphone are the only available digital tools of any power, so I include it here.

Notepad and pencil: Finally – and it usually is my *last* resource tool, if I get that far – is the notepad and pencil. In many projects, my skill at searching and being able to master techniques quickly from a *Wikipedia* description or a *YouTube* video is all I need to be able to work on the problem at hand. On other occasions, however, nothing I find seems quite to fit the problem. I have to adapt a technique that looks like it might work, or I need to re-think my original approach to find another way to bring the power of mathematics to bear on the original problem. This is when, as a mathematician I tend to get very interested in the problem: something new is required! At that stage,

nothing comes close to a paper and pencil – preferably supported by a waste-paper basket – as I play around with the problem looking for new ideas.

Wicked problems

I used the word "problem" quite a lot in the last section. Evidently, I meant "mathematics problem." But what exactly ***is*** a mathematics problem? There are two common interpretations of the term that are both highly relevant here.

First, there are the kinds of procedural challenges you find in a standard mathematics textbook. Questions that ask you to execute a particular calculation to determine a certain value, or to solve a given equation. Mathematics problems of that kind are the standard fayre of school mathematics teaching and surely will remain so. Though there is no longer a need for future citizens to be able to solve such problems quickly and accurately by hand, with real data, the process of learning how to do so with relatively simple data, absent a focus on speed and a high degree of accuracy, remains valuable as a means of understanding the powerful techniques that mathematics provides. Since the solution to those kinds of mathematics problems is now done by machines, that educational benefit is their only value. But it is a significant value.

In contrast to the schools, universities tend to focus on the kinds of "real-world" problems that arise in today's world, for which the various mathematical technologies are not, on their own, sufficient. For an example of such a problem: *Is it a good idea to install solar panels to power your home?* Seems a simple question. Yes or no? But it is far from that.

To answer a question like that most of us begin by comparing several mental images, one of a bank of solar panels on a roof, another of a smoke-emitting, coal-fired, power plant, another of a nuclear power plant, and perhaps one of a wind turbine. We can quickly list pluses and minuses for each one.

Given how aware we are today of the massive dangers of climate change resulting from the emission of greenhouse gases, we probably dismiss the coal-fired power plant right away. But for the other three, you really need to look at some data.

For example, solar panels seem to be clean, and they make no noise, they require very little maintenance, and unlike wind turbines, they don't kill birds. But what is the cost of manufacturing them (including the mining and processing of the materials from which they are made), both monetarily and in terms of impact on the environment? What about the cost of disposing of them when they fail or become too old to function properly? Without some hard data, it's impossible to say whether they are the best choice we might initially see them as.

In fact, as soon as you set aside an hour or so to think about this problem, you start to realize you are being drawn into a seemingly endless series of "What if?" and "What about?" questions, each requiring data before you can begin to try to answer it.

For example, what if a house with a solar-paneled roof is burned in a wildfire, a possibility that residents in many parts of the western United States now face every year? Do those solar panels spew dangerous chemicals into the atmosphere when they burn at very high temperatures? How big a problem would that be? What if, as increasingly happens these days, an entire community burns? How many homes need to burn for the concentration of chemicals released into the atmosphere to constitute a danger to human life?

You are clearly going to have to use mathematics as a tool to collect and analyze the data you need to make some reliable comparisons. So, it's very definitely a mathematics problem. But it's also clear that "doing the math" is the easy part — or rather, the eas**ier** part, especially when there are digital tools available to do all the calculations and execute all the procedures. But what numbers do you collect? Which factors do you consider? Which of them do you decide to include in your comparison dataset and which to ignore?

The point is that school mathematics is a highly constrained domain defined by formal rules. School mathematics problems are unambiguously defined and have single correct answers; moreover, answers that can be obtained within a reasonable length of time. Additionally, each one usually requires just one mathematical technique to solve it. In contrast, real-world mathematics problems tend to be ambiguous, admit to more than one solution, possibly only an approximate solution, and frequently require two or more techniques to solve them.

Since the late 1960s, social scientists have used the term *wicked problem* to refer to ambiguous, open-ended, real-world problems like the solar-power problem above. (Rittel & Webber, 1973) In 2015, the term *kind problem* started to appear to refer to the contrasting, precisely articulated, highly constrained school problems that admit obtainable, unique answers.

Pretty well *any* problem that arises in the social sciences, or in business, or just in life in general, is a wicked problem. The kind problems, the ones the school mathematics focuses on, really exist only in the classroom, where they are used for students to develop mastery of mathematical procedures.

Though my discussion in the previous section of the toolbox that I and other mathematicians use in solving a mathematical problem definitely applies to the solution of kind problems, the ones I was thinking of – the ones I worked on in my many years of applied work – were all wicked problems.

How do we teach mathematical thinking?

Mathematical thinking is really the only game in town when you want to harness the power of mathematics to solve real-world, wicked problems. But how do you teach it? The simple answer is that to help students develop that way of thinking, you have to present them with a series of real-world problems requiring that approach and support them with appropriate coaching. Since that really requires a good grounding in procedural mathematics, together with the ability to learn new mathematical techniques when required, it's not possible to go far down the wicked-problems path in the schools. So, what can a school mathematics teacher do to prepare the way for their students to tackle wicked problems *later in life*? What specific *mathematical* skills should they help their students develop?

As I referred to earlier, it doesn't really matter what mathematical topics are covered, what counts is *how* it is covered. So let's assume it is the topics I mentioned: arithmetic, linear algebra, geometry, and data science.

The goal is to help school students develop a way of thinking. The essence of that thinking is to focus on the concepts, not any particular representation of them. The initial questions a student should ask when faced with a new problem are "What does this mean?" and "Where do I want to end up?" The question "What standard technique(s) do I need to use here?" may come later, or not at all. It's all about understanding the problem.

I'll illustrate the mathematical thinking approach with a simple example. It's one that is often presented as an arithmetical puzzle. For all its seeming simplicity, it captures the spirit of how today's mathematicians work and how their approach is different from both the layperson and from the approach taught in the traditional mathematics class.

Here is the puzzle. (You might want to solve it in your head before you read on.)

> A bat and a ball cost $1.10. The bat costs $1 more than the ball. How much does the ball cost on its own? (There is no special pricing deal.)

The most common layperson's answer is that the ball costs 10¢. That answer is wrong. What leads so many astray is that the problem is carefully worded to run afoul of what under normal circumstances is a successful strategy. As they read through the problem statement, they come to that key phrase "cost more," and they think, "I will need to subtract." They then take note of the data: those two figures $1.10 and $1. So, without hesitation, they subtract $1 from $1.10 (the smaller from the larger, since they know the answer has to be positive). And that gives them the answer 10¢.

Notice, they do not really perform any calculation. They begin by recognizing a linguistic pattern. The numbers are particularly simple ones. Almost certainly, they retrieve from memory the fact that if you take a dollar from a dollar-ten, you are left with 10¢. They might even have visualized those amounts of money in their hand.

Notice too that they clearly understand the mathematical concepts involved. Indeed, that's why the wording of the problem leads them astray!

What they are doing is apply a trick (technically, it's a *heuristic*, of which more in a moment) they have acquired over many financial transactions and most likely a substantial number of arithmetic quiz questions in elementary school. In fact, the timed tests in schools actively encourage such a "pattern recognition" approach. For the simple reason that it is fast and usually works!

We can, therefore, formulate a hypothesis as to why so many people "solve" the problem the way they do. Over time, they had developed a heuristic for problems like this (identify the arithmetic operation involved and then plug in the data) that is (a) fast, (b) requires no effort, and (c) usually works.

In general, a *heuristic* is any approach to problem-solving that employs a practical method that, while not guaranteed to be optimal, perfect, or even rational, is nevertheless sufficient for reaching an immediate, short-term goal or approximation. In any domain, experts, through many years of experience, develop a range of heuristics they have learned are effective.

If a task or problem has a unique, correct solution where the correctness can be verified, the expert will, if possible, carry out the verification. But they will almost certainly *arrive* at the solution by a heuristics path. It's much quicker that way, and it avoids getting lost in a thicket of details where minor errors can cause delays.

Such is the case in pure mathematics, of course. Professional mathematicians invariably arrive at a solution to a complex problem using heuristics, but results are in general (though not always) only accepted by the community after they have been subject to rigorous verification by others.

For real-world, wicked problems, however, such certainty is never possible. In such cases, checks can and usually are made wherever possible, but the answer will inescapably come with some uncertainty, for which it may or may not be possible to give a reliable estimate of error.

For instance, the field of diagnostic medicine is rife with reliance on heuristics, backed up with whatever tests can be carried out in the time available. But as we all know, diagnoses are frequently given in the form, "There is an X percent chance of Y in the next twelve months."

Although it leads you astray from the bat-and-ball problem, the layperson's heuristic approach is a smart one to adopt for the problem, since it uses

something the human brain is remarkably good at – pattern recognition – and avoids something our minds find difficult and requiring effort to master (namely, arithmetic calculation).

Of course, primed by the context in which I presented this particular problem, and most likely being, not a layperson but a mathematics educator, you probably expected there to be a catch. Why else would I give this example? So, after your mind initially jumped to the 10¢ answer (which most of us do), you maybe decided to play it safe and apply a known algorithm. Namely, you likely reasoned as follows (either on paper or in your head):

> Let x = cost of bat and y = cost of the ball. Then, we can translate the problem into symbolic form as: $x + y = 1.10$, $x = y + 1$
>
> Eliminate x from the two equations by algebra, to give: $1.10 - y = y + 1$
>
> Transform this by algebra to give: $0.10 = 2y$
>
> Thus, dividing both sides by 2, you conclude that: $y = 5$¢.

This time, the answer is correct.

Whether or not you can do the calculation in your head, the approach is of course, entirely formulaic and routine. Unlike the first method I looked at (a *heuristic* that is *fast* and *usually right*), this method is an *algorithmic procedure*. As such, it is *slow* (much slower than the first method, even when the algebraic reasoning is carried out in your head), but it *always works*. It is also an approach that can be executed by a machine. True, for such a simple example, it's quicker to do it by hand on the back of an envelope, but as a general rule, it makes no sense to waste the time of a human brain following an algorithmic procedure, not least because, even with simple examples it is easy to make a small error that leads to an incorrect answer.

But there is another way to solve the problem. It's typical of the ways professional mathematicians solve it. Like the first method I gave, it's a heuristic, hence instinctive and fast, but unlike the first heuristic method, it is much harder to fool it with a cunningly worded problem.

This third method requires looking beyond the words, and beyond the symbols in the case of a problem presented symbolically, to the *quantities* represented. Fig.7 more or less captures what the pros do. I should note that I (and likely other mathematicians) don't visualize it quite the way I am about to describe it. In my case, it's more of a vague sense-of-size. The figure I present is the way I explained my method on paper when asked by a cognitive scientist to do so.

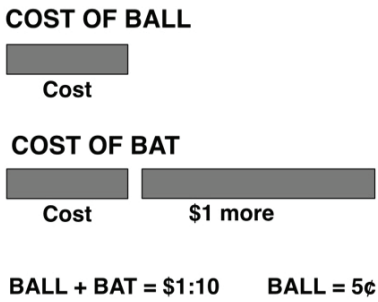

Figure 7. A professional's heuristic solution to the bat-and-ball puzzle.

As we read the problem, we form a mental sense of the two quantities, the cost of the ball-on-its-own and the cost of the bat-plus-ball, together with the stated relation between them, namely that the latter is $1 more than the former.

From that mental image, where we see the $1.10 total consists of three pieces, one of which has amount of $1 and the other two are equal, we simply "read off" the fact that the ball costs 5¢.

There. No calculation, no algorithm. Pure pattern recognition.

This solution is an example of mathematical thinking in action, specifically number sense. It is hard to imagine how a computer system could solve the problem that way. It requires *understanding* the situation described and what is to be determined.

It is still a heuristic. But whereas the layperson's heuristic depends on linguistic patterns, the mathematician's heuristic works with the actual *numerical quantities*, and thus cannot be fooled by tricky wording.

The Australian (pure) mathematician Terrence Tao (2015) has called those three ways of solving the bat-and-ball problem, respectively, *pre-rigorous* thinking, *rigorous* thinking, and *post-rigorous* thinking. Fig. 8 provides a graphical summary of Tao's categorization of the three kinds of mathematical thinking we can bring to problem-solving.

In the blog where he introduced that classification (Tao, 2015), Tao discussed how professional mathematicians solve abstract problems in pure mathematics. The formal, symbolic, rigorous description common in papers and books comes primarily at the end, to check that the solution is logically correct, or at various intermediate points to make those checks along the way. However, the key thinking is post-rigorous, namely an expert's heuristics.

In short, the professional *goes beyond the numerals and the symbols and reasons with the semantic entities those linguistic elements represent*. That is *post-rigorous,* mathematical thinking. And that is what we want to get our students to be able to do.

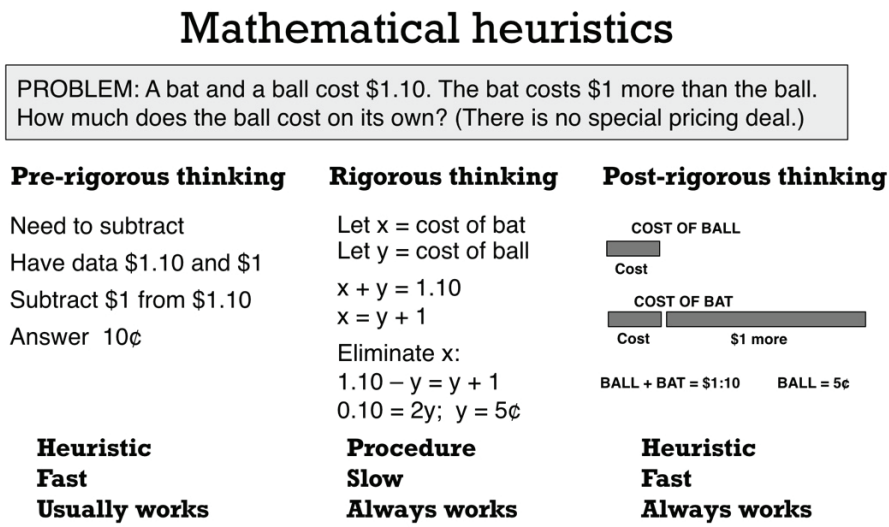

Figure 8. Terrance Tao's three-part classification of mathematical thinking applied to the bat-and-ball problem.

In other words, the question we are addressing in this section can be reformulated as: *How do we best teach students to be good post-rigorous mathematical thinkers?*

Well, we know from research in cognitive science that post-rigorous thinking is almost certainly something that *emerges* from *repeated* practice with rigorous thinking (see, for example, Willingham 2010).

In the days when the only way to acquire the ability to use mathematics to solve real-world problems involved mastering a wide range of algorithmic procedures, professional mathematicians frequently developed into post-rigorous thinkers automatically—as a result of spending thousands of hours doing procedural (i.e., rigorous) mathematics!

But with the range of tools available to us today, there is good reason to assume that, with the right kinds of educational experiences, we can significantly shorten (though almost certainly not eliminate) the learning path from pre-rigorous mathematical thinking, through rigorous, to post-rigorous. The goal is for learners to acquire enough effective heuristics to get them started.

To a considerable extent, those heuristics are not about "doing math" in the traditional sense. Rather, they are focused on making efficient and effective use of the many sources of information available to us today. But before anyone throws away their university-level textbooks, it is important to be aware that the intermediate step of mastering some degree of rigorous thinking is probably essential.

Any increased efficiency in the education process will undoubtedly come from teaching formal methods in a manner *optimized for understanding*, as opposed to optimized for attaining procedural efficiency, as it was in the days when we had to do everything by hand.

Given today's technological toolkit, including search, social media, online resources like *Wolfram Alpha* and *Khan Academy*, and a wide array of online courses, it is surely possible to master most of the rigorous thinking you need "on the job," in the course of working on meaningful, and hence motivational and rewarding, real-world problems.

This is not to say there is no further need for teachers. Far from it. Very few people are able to become good mathematical thinkers on their own. Individuals such as Newton and Ramanujan, both of whom achieved great things with just a few books, are extremely rare. The majority of us need the guidance and feedback of a good teacher.

But, whereas the process of doing mathematics was, until a quarter-century ago, dependent on being able to perform calculations of various kinds, a skillset that the brain does not find natural and which requires considerable training and practice, given the readily accessible calculation tools at our disposal, mathematical praxis today consists largely of using the brain in a manner it finds far more natural: *analogical* reasoning, rather than the logical reasoning previously required. That is the kind of (technology-supported) reasoning our mathematics teaching needs to focus on.

Citations

Abramovich, S (2013). Computers in mathematics education: An introduction. *Computers in the Schools, Interdisciplinary Journal of Practice, Theory, and Applied Research*, *30*(1–2) 4–11. https://doi.org/10.1080/07380569.2013.765305

Devlin, K (2012a). What is mathematical thinking? *Devlin's Angle,* http://devlinsangle.blogspot.com/2012/08/what-is-mathematical-thinking.html.

Devlin, K (2012b). *Introduction to Mathematical Thinking*, Amazon Publications.

Devlin, K (2018). Stanford University Mathematics Outreach Project, *Blue Notepad Videos* https://sumop.org/videos/.

Ferguson, N. *et al.* (2020). Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand, *Imperial College COVID-19 Response Team, 16 March 2020* https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf.

Ma, L. (November 2013). A critique of the structure of U.S. elementary school mathematics. *Notices of the AMS*, *60*(10), 1282–1296. https://doi.org/10.1090/noti1054

NRICH, University of Cambridge (2020). https://nrich.maths.org/mathematically.

Rittel, H., & Webber, M. (1973). *Dilemmas in a general theory of planning. Policy Sciences, 4*(2), 155–169. https://doi.org/10.1007/BF01405730

Singh, P. *et al.* (2018). The use of problem-solving heuristics approach in enhancing stem students development of mathematical thinking. *International Electronic Journal of Mathematics Education*, *13*(3), 289–303. https://doi.org/10.12973/iejme/3921

Stacey, K (2006). What is mathematical thinking and why is it important?. *ResearchGate, Computers in Mathematics Education.* https://www.researchgate.net/publication/254408829_WHAT_IS_MATHEMATICAL_THINKING_AND_WHY_IS_IT_IMPORTANT.

Tao, T (2015). *There's more to mathematics than rigour and proofs*. Aug 20 https://terrytao.wordpress.com/career-advice/theres-more-to-mathematics-than-rigour-and-proofs/.

Willingham, D. (2010). *Why don't students like school? A cognitive scientist answers questions about how the mind works and what it means for the classroom*. Jossey-Bass.