

The Importance of Ethics in Mathematics

Maurice Chiodo (King's College, Cambridge, UK) and Toby Clifton (Fitzwilliam College, Cambridge, UK)

Mathematics is useful because we can find things to do with it. With this utility ethical issues arise relating to how mathematics impacts the world. Now more than ever, we mathematicians need to be aware of these, as our mathematics, and our students, are changing society. In the first of a two-part series on Ethics in Mathematics, we address why, as mathematicians, we need to consider the ethics of what we do.

Mathematics and the world

We study one of the most abstract areas of human knowledge: mathematics, the pursuit of absolute truth. It has unquestionable authority. But, in some sense, absolute truths have absolutely no meaning. The statement “ $2+3=5$ ” is an absolute truth, but what does it mean? Its meaning and utility are added later when people who understand the statement reconcile it with the physical world. It is the mathematically trained who interpret and apply mathematics to the real world and thereby assign it meaning; through this it becomes useful.

Indeed, it is clear that mathematics is one of the most useful and refined tools ever developed. When something is useful, however, it can often also be harmful; this can be either through deliberate misuse or ignorance. The humble knife provides an illustration of the principle; in order to use such a tool responsibly, one must be made aware – often by those who first introduce one to it – of the potential dangers. If a primitive tool like a knife can be so useful and harmful at the same time, then what is mathematics capable of? Mathematics has many more applications, and by the same reasoning must also have a greater potential to do ill. As mathematicians we are seldom warned of this. Other disciplines such as law, medicine and engineering have, for a long time, addressed the potential for harm within their field. We, as the practitioners and wielders of mathematics, need to be similarly aware, and adjust our actions accordingly, otherwise we can, and sometimes do, cause harm with our work. But how could mathematics possibly be harmful, and what exactly might this harm be?

In this article our emphasis is on the experiences of pure mathematicians, although our arguments apply equally to applied mathematicians, statisticians and computer scientists. Many of us (although certainly not all) are motivated to study mathematics by its beauty and intrinsic interest rather than its applications in science and industry. It's as though we are studying a form of abstract art; far from real-world impact or considerations, and only fully appreciated by a small number. Despite all this, government and industry pay for our work; one suspects they don't just do this for the sake of our intellectual stimulation. If our work is completely abstract and detached – an art form, so to speak – then shouldn't we seek funding from those who fund abstract

art? So what might be the value that science councils and industry see in what we do? It is not only the mathematical results that we produce, but also the mathematicians we train. Our mathematics makes a difference, and our students go out and do real things with their training. If it is the case that our work is being funded because it has perceived impact, then surely we should query and understand why we are being paid to do it.

There is already much discussion of ethics in the mathematical community. However, these discussions usually focus exclusively on issues *within* the community. These are important and many of us are already familiar with them: from improving diversity and inclusivity, to widening participation in mathematics, to addressing instances of plagiarism and publishing irregularities. These are pressing concerns. Every discipline engages with such *intrinsic* ethical issues. These, however, are not our focus in this article. Mathematics is one of the few disciplines that fails to address *extrinsic* ethical issues; those concerning how the community impacts wider society. This particularly includes ethical implications relating to the applications of mathematics and work of mathematicians. It is these extrinsic ethical issues that we are trying to raise awareness of. Our concern is not so much that mathematicians are deliberately malign, but instead that they fail to recognise these extrinsic ethical issues. Indeed, most of the mathematicians we have come across would balk at the thought of acting unethically; the problem, instead, is that many do not recognise that mathematical work *can* have such an effect.

Some case studies

Having recognised that mathematics is useful because it can be applied, and that with these applications come extrinsic ethical issues, we now consider two concrete examples: the global financial crisis (GFC) of 2007–8 and targeted advertising.

The GFC was one of the defining events that shaped the modern global economy. Its repercussions have been felt around the world, with many suffering a decline in living standards. The causes of the GFC are complex; however, there is consensus that mathematical work played a vital role. An important factor is thought to have been the misuse of Collateralised Debt Obligations (CDOs). These saw mathematicians pool large collections of interest-bearing assets (mostly mortgages), then ‘cut the pool into pieces’ to form a collection of interest-bearing products. Mathematically these products had less overall risk and thus higher value than the original assets. They were traded wildly. The mathematics behind their construction is highly non-trivial, requiring stochastic calculus, differential equations, etc. Research mathematicians, beginning with the work of Black and Scholes, and later Li, derived a model and pricing for-

mula for CDOs. Though it took a deep understanding of mathematics to derive these models, only a more superficial understanding (at the undergraduate level) was required to apply and to trade them. As a result, their users may not have fully appreciated their limitations or inner workings. Mathematics – which by itself is sure and certain – seemed to explain their value, and so most were happy. Unfortunately, some of the assumptions did not hold. For example, the model assumed there wasn't tail dependence in the default risk of underlying assets, but there was; for instance when two mortgaged houses were on the same street. In the end, the risk was not properly accounted for, and when house prices declined it led to the write-down of \$700 billion of CDO value from 2007 to 2008. The rest is history.

Our next example is targeted advertising. Adverts have always been placed so as to catch the eye of their desired audience. However, now that people possess portable internet-connected devices and social media accounts, it has become possible to target adverts at the individual level. Nowadays, these can be tuned to fit very specific demographics, and as such it's now possible to specify who *doesn't* see an advert. This allows advertising campaigns that are selective, contain adverts that contradict each other, and that are impossible to externally scrutinise. In short, adverts can now be used to manipulate individual people. This becomes particularly dangerous when applied to political advertising. Using large data sets obtained through social media it is possible to profile the political persuasions and preferences of an individual. Machine learning has become the main tool of the trade here, and it is the mathematically trained doing it [3]. These adverts can even deceive by appearing non-partisan. For instance, one can send an advert saying "Voting is important; make sure you vote" only to those who might be inclined to vote for your party.

Whatever the strategy, these types of adverts are increasingly prevalent, and it is thought that such tactics influenced the 2016 US election and UK referendum on EU membership. It is we mathematicians who make all this possible. Cambridge Analytica, one of the organisations alleged to have been involved in such advertising, had a small team of no more than 100 data scientists [4], some of whom were trained mathematicians. Regardless of one's political persuasion, it is clear that this sort of work is deceptive and dangerous, and that mathematicians are enabling it. Ultimately, it is mathematicians who make up part of the teams specifying how such targeting works and carrying it out.

The impact of mathematicians

As a result of the pace and scale at which modern technology operates, through use of internet connectivity and readily-available fast computation, the consequences of the actions of mathematicians are more quickly realised and far-reaching than ever before. A mathematician in a big tech company can modify an algorithm, and then have it deployed almost immediately over a user base of possibly billions of people. Even on a smaller scale, we have seen that a small number of mathematicians,

despite limited resources, can have a vast impact on the world; targeted advertising exemplifies this.

If you model a physical system, such as gravity, then your model is falsifiable. If the model does not accurately reflect the physical system, then on application it clearly fails – your rocket doesn't launch properly. You know when a model was good because the rocket makes it to the moon and back. Modelling a financial system is more difficult, as the system is affected by the application of the model. A pricing algorithm, if widely used to buy or sell a product, influences the market for the product in question. How does a model model its own impact?

So now what happens if you are modelling the future behaviour of people by predicting something like: 'How likely is a particular individual charged with a crime to reoffend with a serious offence, a non-serious offence, or not reoffend, in the next 24 months?' Furthermore, what if that is being used to determine what prosecution and sentencing mechanisms are applied to that person?¹ If you predict that a person will reoffend seriously in 24 months, and they don't (after being released or acquitted), then you might observe that. But what if they are found guilty and sentenced to 25 months with the choice of judicial process based on your prediction? How do you test whether your prediction was correct? Now we have a serious ethical issue: we are using mathematical reasoning to make decisions about people that impact their lives, and in many of these cases we can never know whether the decisions made were desirable or appropriate. Is it right to use mathematics in such a way without careful reflection?

We now face an ethical dilemma. Do we limit ourselves to falsifiable claims, or do we allow ourselves to make claims, make decisions and initiate actions that are unfalsifiable? We are of course entitled to do the latter; however, we should then bear in mind that we have lost mathematical certainty. Furthermore, if we do this, we should broaden our perspective and training so that we can incorporate as many aspects of society as possible.

Concerns for the future

So what is on the horizon for mathematicians? Is it sufficient to simply look at the above list of cases and avoid those specific actions or industries entirely? Unfortunately not; new mathematics produces new ethical issues every day. Such a future example may lie in alternative credit scoring. This is starting to be done by new companies who lack access to standard datasets that established credit-scoring agencies have (such as financial records, bill payment history, etc.). They instead use different datasets such as social media profiles, in some cases requesting full access to social media accounts by asking

¹ The Harm Assessment Risk Tool (HART) developed by the Durham Constabulary, which uses random forest machine learning, is used to make such predictions, and then determine if an accused criminal is to be offered the opportunity of going through the *Checkpoint* program (tinyurl.com/y4vxrd77) which is an alternative to criminal prosecution aimed at reducing re-offending.

for login credentials [2]. While this sounds undesirable to the point that most people will not be interested, it must be remembered that some people will be sufficiently desperate for credit so that there will always be some takers. These companies scrape an applicant's social media looking for actions they perceive to reflect creditworthiness. These could include places the person visits, the hours they sleep, the 'quality' of the friends they have and so on. This approach is unfalsifiable, lacks proper regulation and has the potential to harm society since the extension of credit is a mechanism of social mobility. If such a process, one that is enabled by mathematically trained people, starts having negative impact, who is accountable? Ultimately, we must live in the world that we and our students create, and we must ponder whether there is a sense in which we are partially responsible.

Do these ethical issues arise in academia?

But what about mathematicians working in academia: are any of these ethical issues relevant to them? Consider a pure mathematician, a number theorist, say. Suppose they develop an algorithm for fast factorisation. Should they publish it? If so, when, where and how? If not, what should they do? Should they have thought about it beforehand? We have asked many mathematicians this exact question, and a typical response is: "I would publish it on arXiv immediately. It's my right to publish whatever mathematical work I do." (Not all mathematicians give such a response, but many do.) When pressed on the consequences of publishing such an algorithm in that way – for instance the breaking of RSA encryption in a chaotic manner and the ensuing collapse of internet commerce and the global economy that would follow – one explained: "Well, it's their fault for using RSA. It's not my problem." Of course, responsible disclosure is a complicated topic, and one that is heavily debated by security researchers. But with an example like this, ethics has crept into the world of the pure mathematics researcher in academia.

If an area as abstract as number theory is not 'safe' from ethical considerations, is there any mathematical work that is? Can a pure mathematician hide from ethical issues in academia? What about a statistician, or an applied mathematician? Or do ethical questions arise for all mathematicians regardless of where we do our mathematical work?

Why management can't guide us

Some mathematicians (academic or industrial) may think that, since they are not directly involved in the application of their work, they need not consider its extrinsic ethical implications. After all, we just do the maths, and so it's 'not our problem'. This oft-held belief is generally associated with the perception that there are people and structures above us (managers, supervisors, advisory boards, etc.) who will intervene to prevent us from doing anything that we ought not to. *We* work on the abstract problems, *they* worry about why. But can we rely on management to do this effectively? Will they vet our work to ensure that its use is aligned with the values

of society? At each stage of separation from mathematical work some understanding of it is lost. It is difficult for a manager to understand *all* of the mathematical work we do and its limitations when applied and used. It is the nature of management that managers will only have partial knowledge of the work being done. There would be no point in a manager reproducing the work of all of the people under them, and mathematics is such that if you don't 'do it for yourself' then there is a chance you may not fully understand it. Given this fact, there is always an onus on the individual mathematician to consider the ethical implications of what is being done. Of course, it must also be considered that managers might have other values, perhaps more aligned with the objectives of the organisation than of wider society. We should understand and anticipate this.

Some managers may go so far as to try to manipulate us. For instance, if we voice objection at what we have been asked to do, they may try to quash it with the classic argument: "If you don't do it, then someone else will". At a first glance, this seems convincing; however, it fails on two counts when referring to mathematical work. First, there are not that many mathematicians in the world. We possess a unique set of skills and abilities, and it requires years of training to produce a good mathematician, even when starting with someone who has the right interests and reflexes. Given the scarcity of mathematicians, this argument fails in practice. Moreover, as mathematicians we understand its contrapositive; the original statement is equivalent to: "If no one else does it, then you will". This is, of course, absurd. The argument has as the implicit underlying assumption that the task being requested will definitely be completed. If no one else builds me a nuclear bomb, then will you? What we should really be considering here is the argument: "If you don't do it, then someone else *might*". True, someone else might, but they may not be easy to find, or even exist at all. Now the power of meaningful objection has returned to the mathematician. Whether you choose to take the pragmatic perspective that there are not many mathematicians or the logical perspective of the contrapositive, your objection means something. Some mathematicians take this idea even further, and take a conscious decision to take a seat at the table of power, effecting positive change from the managerial level. This happens in various areas: in academia, in industry and even in politics. This is discussed in more detail in [1] as 'the third level of ethical engagement'.

Why the law can't guide us

The problem extends beyond management. We may think that the law provides a clear description of what is and is not acceptable to society, and thereby presumably what is and is not ethical. However, this misses the point for several reasons. Firstly, the law is not an axiomatised system; it is interpreted by courts rather than by machines. This is a type of system with whose details mathematicians are generally not familiar. Furthermore, there is the problem that the law will always lag behind technological development; we cannot expect lawmakers

to have done our mathematics before we do it ourselves. Additionally, the processes by which laws are made are (deliberately) slow, requiring public consultation, votes and implementation periods. Consider the case of the General Data Protection Regulation (GDPR). It started to be written in 2011, only came in to effect in 2018, and is thought by many to be already out of date. Finally, it can be the case that lawmakers lack a full understanding of the fine details of the subject at hand. For example, a member of the UK Science and Technology Select Committee, Stephen Metcalfe, declared at a public outreach event that “one solution to algorithmic bias is the use of algorithms to check algorithms, and the use of algorithms to check training data”. Ultimately, the law is not there to serve as moral advice; there are plenty of immoral things one can do that do not break any laws. As such, it is not well suited as a source of ethical advice.

Thus, if we can't rely on management and we can't rely on lawmakers and regulators, then who can we rely on? The answer is as obvious as it is difficult to admit: ourselves. The only way mathematicians can try to prevent their work from being used to do harm is if they think about it themselves. No one else can, so we must.

A growing awareness of Ethics in Mathematics

Awareness that mathematicians need to consider extrinsic ethical issues is building in the community. In 2018 the head of mathematics at Oxford, Professor Michael Giles, commented at a panel discussion event: “Cambridge Analytica is interesting from one point of view in that, if you'd asked me 20 years ago whether mathematicians at the PhD level needed to be exposed to ideas of ethics, I would have said ‘Clearly, that is irrelevant to mathematicians’. Now I really think that this is something we have to think about. In the same way that engineers have courses looking at ‘What it means to be a professional engineer’, and ‘Ethics, and your responsibilities as an engineer’, I think that is something that we have to think about as mathematicians now.” Moreover, arxiv.org is currently revising the description of their mathematics tag *History and Overview* to include “Ethics in Mathematics” as a sub-category.

As part of their formal training, few mathematicians have ever been told about extrinsic ethics before. Previous generations of mathematicians have evaded this crucial point, and in the process have possibly let society down. It rests on the current and upcoming generations to pick up this idea before it's too late.

Mathematicians always take a generation or more to accept a new and fundamental idea about the nature of their subject; debates about the admissibility of zero as a number provide such an example. We're at a similar juncture again. Now some say: “Surely there's no use in considering ethical issues in mathematics”, but by the time our students are professors and industry leaders, they may well be saying: “Of course we should be considering ethical issues in mathematics!” But why hasn't the mathematical community taken this on board already? Why wasn't this done 100 years ago, by the likes of Gödel and Russell? Two reasons come to mind. Firstly, the dan-

gers were less proximate, since much of today's technology simply did not exist. Secondly, every mathematics undergraduate was already exposed to philosophy, as it formed part of every university education. Thus, exposure to Ethics in Mathematics, in its own right, was less urgently needed.

So if Ethics in Mathematics has become so important to mathematicians, then how might we teach it to them in a relevant and useful way without foisting an entire philosophy degree upon them? Disciplines such as law, medicine and engineering have long taught their undergraduate students about extrinsic ethics in their respective fields. In the following article of this *Newsletter* we'll further explore why such teaching has not yet occurred in mathematics, and outline how one might go about giving such directed teaching of Ethics in Mathematics (EiM).

Acknowledgements

We wish to thank Piers Bursill-Hall and Dennis Müller for their valuable discussions over the past few years, and Dennis in particular for his comments and suggestions on previous versions of this article. We also wish to thank Iain Moffatt, as editor of the LMS Newsletter, for his thoughtful contributions on revisions of this article.

This article first appeared in the *Newsletter of the London Mathematical Society*, Issue 484, September 2019.

References

- 1 M. Chiodo, P. Bursill-Hall, Four levels of ethical engagement, *Ethics in Mathematics Discussion Papers*, 2018/1 (2018).
- 2 Systems and methods for using online social footprint for affecting lending performance and credit scoring, Patent No. US 8,694,401 B2, 8 April 2014.
- 3 Cambridge Analytica: how did it turn clicks into votes?, *theguardian.com*, <https://www.theguardian.com>, tinyurl.com/yy89e5v1, May 6th 2018.
- 4 Secrets of Silicon Valley, Series 1 Episode 2: The Persuasion Machine, see recording at 30:17. BBC documentary, 2017



Maurice Chiodo [mcc56@cam.ac.uk] is a Fellow and Teaching Officer at King's College, Cambridge. He is the lead investigator of the Cambridge University Ethics in Mathematics Project (ethics.maths.cam.ac.uk), developing a programme to teach mathematicians about the ethical implications of their work.



Toby Clifton [cmctc3@cam.ac.uk] is a recent graduate of astrophysics, and is the current president of the Cambridge University Ethics in Mathematics Society (<https://cueims.soc.srcf.net/>).