An Introduction to Mathematical Research for Undergraduates

by Matthew Romney

1. What is mathematical research?

Mathematical research is the process of creating new mathematical knowledge. In this sense, the word *research* might be something of a misnomer—after all, an artist or musician doesn't refer to their creative work as "research". Nevertheless, the term has become standard, probably because it reflects the professionalization of mathematics as an academic discipline. In common discourse, *research* refers to studying a topic, finding out what is known about it, and organizing this information in a useful way. Mathematical research includes some of this but it is much more. It is a creative activity and usually a collaborative one.

Doing research commonly takes the form of *solving a problem*. To frame a mathematical topic as a problem makes it concrete. This allows you to focus your mental energy and measure your progress. At the community level, problems are an efficient way to communicate mathematics—and allow the eventual solver to get his or her due credit. Indeed, a fundamental problem that remains unsolved for a long time will acquire a certain fame and prestige, all of which can be claimed by the person who solves it. The great mathematician David Hilbert famously suggested 23 problems in his 1900 International Congress of Mathematicians address, which went on to guide research into the following century. In this address, Hilbert poetically says: "A mathematical problem should be difficult in order to entice us, yet not completely inaccessible, lest it mock at our efforts. It should be to us a guide post on the mazy paths to hidden truths, and ultimately a reminder of our pleasure in the successful solution."

Attempting to solve a famous problem is risky business, and certainly not recommended for a beginning researcher. But the same joy and thrill of problem solving is available to anyone attempting to do mathematical research. One of the beautiful parts of mathematics is its variety in subject matter, style, and difficulty. There is a problem suited for everybody.

2. Why do research?

There are many good reasons to do mathematical research. Here are what I consider to be the main ones.

–*Intrinsic knowledge.* The most obvious reason to do research is for the knowledge produced. Given that humans as a species have a surplus of time beyond simply surviving, we need to do something interesting and meaningful with our existence. What better way than the pursuit of an understanding of the world around us?

To do mathematical research is to join a conversation going back for decades, even centuries or millennia. There is a certain thrill of contributing something new to this conversation, of uncovering some piece of truth about the mathematical universe that has never been noticed before. The abstract nature of pure mathematics does, to a certain extent, separate it from the "real world", but what is gained is permanence.

Your work will be there forever, in principle at least. The mathematical discoveries of the ancient Greeks and other civilizations are still remembered today.

–*Applications.* While most of this article is geared towards pure mathematics, it's also the case that applied mathematics is a large and thriving industry. Eugene Wigner famously spoke of the "unreasonable effectiveness" of mathematics in understanding the natural world. Most of pure mathematics will not foreseeably have real-world applications, but there have been intriguing surprises in recent history (cryptography being an oft-cited example), and the distinction between pure and applied mathematics has to a large degree faded. For example, the most frequently cited work of Terry Tao, arguably the foremost pure mathematician of our generation, is in signal theory, a branch of applied mathematics.

–*Social function.* On a more personal level is the prospect of building a career as a mathematician. Research is the main proving ground for aspiring academics: a chance to show your abilities and build your reputation. Producing good research gives you social capital, so that your words and actions will be taken more seriously by the academic community. You'll also build deep professional and personal relationships with others through research. These are often lasting relationships spanning national, language and other boundaries.

–*Training.* Most research papers are not going to revolutionize the subject of mathematics. In fact, most papers are of immediate interest only to a small group of fellow specialists. However, the experience gained by studying these minor problems often leads to more significant future discoveries. Per the 80/20 (or Pareto) Rule, this is the nature of things: a small share of time or effort accounts for the majority of progress or accomplishment.

Beyond a specific mathematical specialty, there are certain thought processes and skills that are obtained from doing research that seem impossible to get otherwise. Indeed, there is a mentality of a researcher that I believe will have great value wherever you end up in your career. These skills include: the ability to understand others' technical ideas and fill in gaps as needed; the ability to reason heuristically; the confidence to pursue your own problem-solving strategies; the ability to modify or reformulate a problem or connect two seemingly unrelated topics.

–*Interest.* I believe that the most immediate and personal motivation for most mathematicians is that math is interesting. Day after day, this is what keeps us going. In fact, I would even say that if what you're studying isn't interesting, then you're approaching it wrong. You might argue that one of the basic principles justifying pure mathematics is that pursuing what is interesting also leads to what is useful.

3. How do you do research?

The research process can be divided into four basic stages.

A. **Ask questions.** In order to do research, first you need something to study. As already discussed, this is often best framed as a mathematical problem or question. This makes the topic concrete, allowing you to focus on it and measure your progress. Keep in mind that most questions will probably not lead to anything interesting, so it helps to have a large pool to draw from. Don't get stuck for too long on any single problem.

What about finding the subject matter for your questions in the first place? You might ask: what is something not well understood by mathematicians that *should* be? What do you personally want to understand better?

B. **Develop a strategy or method to study your question.** This is usually the most difficult and prestigious part of research, where you can let all of your genius and creativity shine. You'll often be mistaken and many ideas will lead nowhere, but with some time you should also score some hits. Once you have a workable strategy, the rest of the research process should be routine or straightforward.

   That said, you shouldn't focus too obsessively on a given problem. Instead, let the problem guide you to a deeper overall understanding of the topic. This will help ensure that the time you spend is productive, independently of whether you solve your original question or not. Moreover, that deeper understanding is what may eventually lead you full circle back to solving the original problem.

C. **Work out the precise details of your solution.** This is what is being trained by the many homework exercises you do as a student. It can be tedious and time-consuming, but it's very essential to the integrity of mathematics as a discipline and the ability to communicate ideas to others.

D. **Write up your work into a research paper.** The final part is to put all your ideas into an attractive final product. This step is somewhat less important than the others, in the sense that a correct mathematics paper will still probably get published even if it's poorly written. In fact, my experience is that most math papers have significant issues with the quality of writing. But, you want your paper to actually be read and understood, don't you? Then make it easy for the reader.

As an undergraduate researcher, your main job will be items C and D, while a large part of the burden for items A and B will go to your mentor. After all, the hope is for you to finish a completed project in a matter of a few months. However, don't hesitate to try your hand at the first two.

As you progress as a researcher, it will eventually become important to become proficient with all the stages of research. At the very least, you'll need to build up a robust enough network of collaborators to make up for any weaknesses of your own.

4. Stages of your mathematical training

–*Undergraduate studies:* At this point, your main concern is obtaining a strong foundation in preparation for graduate school, or whatever you choose to do in the future. The main component of this is your coursework. Make sure to perform well and be responsible students; the professors who will write letters of recommendation for you will be objective with their assessments of your abilities compared to your peers. Participate and be visible; asking relevant questions in class and attending office hours are good ways to accomplish this.

As an undergraduate, it is common and helpful to try your hand at a research project. This is often done through a Research Experience for Undergraduates (REU) program such as the one being sponsored here at Stony Brook. Such a program can be as useful (or non-useful) as you and your mentor make it. To make an REU as useful as possible, I would recommend working in relatively small groups—this gives the benefits of collaboration while allowing each group member to contribute something meaningful. For example, if the fruit of your REU is a simple paper with, say, seven or more co-authors, you won't reap much career benefit from it, since there is no real way to distinguish your contribution. But papers with two or three, or possibly four, authors are great since it's more evident that everyone contributed. Also, remember that an REU project will usually be fairly modest and elementary in the scheme of things, so don't get so absorbed in your project that it detracts from your overall mathematical education.

During your undergraduate years, and throughout your graduate studies as well, seek out opportunities for internships and other extracurricular programs. There is a good chance you might ultimately choose a non-academic career path, or be forced to fall back on a Plan B if your academic career doesn't pan out as expected. Even if you do pursue an academic career, having outside experience and connections can be a strong asset for several reasons. It can help you advise a variety of students and connect them with future employment, lead to potential multidisciplinary collaborations, improve your skill set and hireability, and provide a different perspective that can influence your academic work.

–*Graduate studies:* The next stage is your graduate studies. Naturally, this requires applying for and getting accepted to graduate school. There are graduate program rankings published by *U.S. News* and others that are a fairly reliable guide. In general, it helps to go to a more prestigious program, though prestige is just one consideration. My experience is that it is quite possible to have a successful academic career from any top 40 (or so) U.S. program. If you are able to get into a top ten program (per *U.S. News*, these would be a subset of Princeton, Harvard, MIT, Stanford, Berkeley, Chicago, Columbia, UCLA, Caltech, NYU, Yale, Michigan), then your prospects improve quite a bit, especially for getting hired at an elite institution. On the other hand, the universities in the 11-40 range are probably roughly equivalent, and other factors probably matter more than prestige of the department itself. These departments include most of the major public universities (Texas, Wisconsin, Stony Brook, etc.), and you really can't go wrong with any of these. Once you get outside the top 40-50 programs, your academic prospects seem to decline significantly. (An exception is if your doctoral advisor is particularly famous; lower ranked departments still often have a distinguished professor among their ranks). What this prestige often means in practice is that you can still land a good postdoc position even with an "average" thesis. After your postdoc, though, it's your own work that counts.

One thing I would advise is to go into graduate school with an open mind. It's fine to have a broad idea of what you'll study, such as geometry or partial differential equations or number theory. But be flexible about the specific subfield until after you have a good working relationship with your thesis advisor. If you've done undergraduate research, don't get too attached to your topic. Instead, be ready to expand your interests towards higher-level fruit. Your graduate studies are a time to immerse yourself deeply in a topic without any expectation to produce new results right away. This is a luxury you likely won't have again in your career, so make sure to take advantage. Quality counts more than quantity when it comes to research, so find a rich and deep topic to become an expert on. Don't rush through your degree, but try to maximize each step along the way.

–*Postdoctoral stage:* towards independent researcher. If your goal is to be a research mathematician, your Ph.D. diploma will be followed by a postdoctoral stage usually lasting about 3-5 years in the United States, after which you will hopefully be a contender for a permanent position. There's really no fundamental difference between what you do now and what you did as a graduate salary (though your salary has fortunately jumped quite a bit and will probably jump some more once you get a permanent job)—except that you'll be expected to do things independently rather than under the supervision of a thesis advisor. Fortunately, if you've done things right, you'll have plenty of colleagues to talk to that will help lead to new projects and collaborations, so "independent" is far different than "isolated". You also may have an official postdoctoral mentor.

Throughout this entire process, it's good to remember to keep non-academic options in mind. Given the basic statistics of new PhDs versus university job openings, it may even be necessary to switch career paths. My impression is that about one third of new PhDs (from, say, a top 30 U.S. graduate program) will leave academia immediately for careers in finance, tech, government or government-sponsored laboratories, and the like. (The American Mathematical Society keeps these sorts of statistics, so please confirm things there. Many graduate programs also list the first jobs of their graduates.) Most of the rest will take a postdoc, with a select few immediately landing a permanent teaching-oriented position. My impression of the current situation is that most new PhDs (again from, say, a top 30 U.S. graduate program) who want a postdoc and have produced an adequate thesis will get a reasonably good postdoc.

The pressure of the academic job market will be felt most at the conclusion of your postdoctoral period, since the supply of desirable permanent positions is fairly tight. You may have to bounce around multiple postdocs. One major way to separate yourself from the pack is to receive a postdoctoral fellowship or grant such as from the National Science Foundation. Make sure to apply to these. That said, the tightness of the job market is a function of your expectations. If you are hoping for a professorship at one of the major research institutions, you will most likely need consistent research output, including publications in the very top journals (see the discussion of math journals below), along with a good teaching record and good visibility in the math community. These jobs are hard to come by, so you might set a more modest goal, such as getting hired at a lower-ranked department.

For any kind of research position (for example, any university with a graduate program), you'll likely most need publications in very good journals. Beyond that, there are a variety of teaching-focused positions. Pay attention to your fellow graduate students who are a few years older than you to see where they end up; that should give some guide of what to expect. Of people who finish a PhD (from, say, a top 30 U.S. graduate program), I would guess about half will find a permanent home in academia; out of these, about a quarter will land a research-oriented professorship, with the rest taking teaching-oriented positions. In summary, the journey ahead will not be easy, but you can have some confidence that things should work out in a satisfactory way.

5. What is a research program?

A research program is an agenda of ideas, problems, and methods that you plan to work on in the future. Since it's hard to predict whether an open problem will get solved or not, the idea of a research program is somewhat paradoxical. As Albert Einstein said, “If we knew what it is we were doing, it would not be called research. Would it?” Nevertheless, when it comes to getting grants, fellowships, and job offers, you need to convince others that your research is important and will likely lead to successful outcomes. Part of this is building up your credibility by having a track record of success. Another part is finding the right balance between the mundane, predictable research projects and the ambitious ones.

You might compare the job of a researcher to a gold prospector from the not-too-distant past. Most land does not have gold, so you might have to travel a bit at first. But once you find something, then don’t just keep on walking. Where some gold is, there’s bound to be more. Take advantage! Indeed, you’ll often find that one successful research project very naturally leads to another, which in turn leads to another, and so forth. This is exactly how mathematical careers usually come about.

You won’t be expected to have your own official research program until the end of your graduate studies when you are applying for postdocs. However, there are opportunities in the meantime if you can get ahead of the curve. For one thing, your graduate school application may fare better if you are able to articulate convincingly what you hope to do as a graduate student. Moreover, applying to programs such as the NSF graduate research fellowship requires you to have the beginnings of a research program even as an undergraduate.

6. How to ask mathematical questions?

Asking good mathematical questions is not easy, and in many cases it may be the most difficult part of doing research. In fact, the rest of the research process can become surprisingly easy once you’ve found the right question to work on. You’ll find that many prospective questions are too simplistic, too difficult, or too easy.

Still, here is a list of basic strategies to keep in mind. Note that these strategies overlap. As you go through your regular coursework, you might get in the habit of asking these sorts of questions—though keep in mind that answering a given question in a complete or meaningful way may be too much work to be practical.

–*Generalization:* Suppose that Theorem A is proved under a certain assumption, or applying to a certain class of objects. Then the natural next step is to try to remove or weaken this assumption, or to broaden the class of objects under consideration.
–*Counterexample:* Very often, Theorem A does not generalize in any clear way, and so your task is to find a counterexample once the assumption in question is removed.
–*Sharpness:* An inequality containing some constant is sharp if the constant is the best one possible. On a more general level, it is always interesting to find the strongest possible form of a theorem. This can include finding the object or configuration that realizes the worst possible behavior—a so-called *extremal problem*.
–*Make quantitative*: Many mathematical theorems are of a qualitative nature. However, it is often useful to have quantitative versions of results: that any parameters present in the solution can be estimated or bounded in terms of suitable parameters in the initial objects.

–*Algorithm*: It is common to have an existence theorem that guarantees solutions to a problem exist. An algorithm is an explicit list of steps to actually find these solutions. Once you have some algorithm, you can try to analyze its efficiency and determine to what extent this can be improved.
–*Inversion:* Suppose that you compute a certain quantity or output for a given input object. You might ask: given knowledge of the output, can you determine what the original input is (or could be)? There is a large branch of mathematics called *inverse problems* dealing exactly with these sorts of questions.
–*Classification:* Your task is to enumerate all possible objects of a given type up to some notion of equivalence. The most famous example may be the classification of finite simple groups, essentially worked out from the 1950s to 1980s. This is a great method to base a research program on: it's usually easy to communicate to others and convince them that it's interesting, and it provides a reliable framework to generate consistent research output.
–*Add an assumption:* If a problem seems intractable as stated, then you can try to solve it under an additional assumption. This assumption may turn out to have independent interest and become its own object of study.
–*Application:* Once you find a method that works for one problem, then look for other problems where the same method can apply.

7. Writing a research paper

There are many good resources on the internet on how to write a mathematics research paper. I will try not to repeat too much of what is already said but instead highlight a few important considerations. Also, perhaps the best way to learn how to write a good research paper is by reading good research papers from other authors.

–*Clarity.* This is the number one rule of writing. Be precise with your definitions. Be precise with your theorems. Be precise with your proofs. Include full details unless the details are literally redundant. Write in such a way that any patient reader with a standard mathematical background can follow your work. You might be concerned that providing too much detail will bore readers, insult their intelligence, or make your paper excessively long. These concerns are almost never well founded. On the contrary, the reader will appreciate your effort and thoughtfulness. They have not thought long and hard about the subject of your work the way you have, and they shouldn't have to think as long and hard as you did in order to read your paper.

Write so that the logic of your proofs is fully transparent. The flow should be, in essence, "Because A, then B. Because B, then C. Because C, then D. …". In fact, there's a sort of poetic quality to this style of writing that communicates the beauty of math, and makes it easy for the reader to follow.

If you're doing research at the level of someone like Mikhail Gromov or Grigori Perelman, then others will make the effort to understand it regardless of how it is written. For the rest of us, we need to write clearly in order to have an audience and for others to take us seriously. The typical research mathematician writes about two papers a year. An extra week or two of revisions is not going to set you back noticeably. Resist the very understandable temptation to post your draft as soon as you can.

–*Motivation.* All math papers look superficially the same. They contain theorems, definitions and proofs, in which some amount of technical work is carried out. Unfortunately, this can make it hard for people other than the authors and a small number of specialists to evaluate the originality and importance of an individual work. One of the parables of Jesus concerns the wheat and the tares; the two plants look indistinguishable despite the latter being a weed. It is similarly hard to distinguish between "good" and "mediocre" math. (In principle, the reputation of the journal that a paper is published in is the mechanism for establishing the quality of the work, but there are too many variables at play for this to be very reliable, especially for the mid-tier journals.)

What all this means for you as an author is that you need to make an effort to present your work well, lest the wheat you have worked so hard for be mistaken for tares. Be explicit about communicating the value of your paper to your reader. Most mathematicians will probably not understand in precise detail the ideas in your work. There are too many papers and too many subfields for this to be feasible. However, a reader should not be left wondering: why is the author interested in this topic, and why are the results novel or significant?

–*Write a good introduction.* Here is a basic checklist to help implement the previous two points. To the extent you are able, do the following: (1) introduce your topic in a non-technical way that any mathematician in the broader field can understand and appreciate; (2) state a clear theorem as soon as possible, postponing overly technical definitions until later; (3) describe any applications or consequences of your results; (4) relate your work to prior work on the topic; (5) outline the ideas and methods of your paper, making it clear what the main novelties are.

8. Submitting your work for publication

Once you're finally done preparing your article (and this always takes much longer than expected in my experience), it's time to submit to a journal for publication. Over time, you'll become acquainted with the standard slate of mathematics journals. There are five journals commonly recognized as the very best pure mathematics journals: *Acta Mathematica*, *Annals of Mathematics*, *Inventiones Mathematicae*, *Journal of the American Mathematical Society*, and *Publications Mathématiques de l'IHÉS*. All these journals are basically interchangeable in prestige (except for possibly *Inventiones*), and it's common for the leading mathematicians to spread their best work among all of them. There are a number of other top generalist math journals (including *Communications in Pure and Applied Mathematics*, *Duke Mathematical Journal*, and *Journal of the European Mathematical Society*) as well as the top specialty journals (such as *Geometric and Functional Analysis*, *Geometry and Topology*, and *Journal of Differential Geometry*). There are also the top applied mathematics or multidisciplinary journals such as those published by SIAM (Society for Industrial and Applied Mathematics) and IEEE (Institute of Electrical and Electronics Engineers), but I am less familiar with these.

Most likely, you won't be submitting your first work to one of these journals, though you should aspire to reach that level eventually. You'll find that there are many options out there, so it may take some thought to find the right journal. There are a number of subject-specific specialty journals that are often considered the best dedicated to that subject (*Calculus of Variations and Partial Differential Equations*, *Ergodic Theory and Dynamical Systems*, *Journal of Differential Equations*, *Journal of Functional*

Analysis, and so forth). These may be a good home for solid, but more everyday, research such as that produced for a graduate thesis. There are also plenty of mid-tier generalist journals, typically sponsored by an academic society (e.g., *Proceedings of the American Mathematical Society*) or a university (e.g., *Illinois Mathematics Journal*), and no shortage of "basic" journals with a lower standard for acceptance. Finally, there are some dedicated undergraduate journals (e.g., *Involve*) and expository journals (e.g., *American Mathematical Monthly*), but I would suggest looking first to publish in a standard research journal.

If you're not sure where to try submitting, there are a few things to try. First (and obviously), you can look for journals specializing in the topic of your paper, and then try to match the prestige of the journal to the quality of your results. Another strategy is to look at the papers that you cite. For example, if your result builds directly on a previous paper, it may be appropriate to submit your paper to the same journal. You can also look at the editorial boards of possible journals to see if there is a familiar name who is more likely to appreciate your work. That said, my experience is that the lack of an editor in your immediate field should not deter you from submitting to a journal. Editors are used to handling a wide variety of papers.

For many mathematicians, another consideration is to submit your work to journals that follow good publishing practices, such as charging a reasonable cost for subscriptions. For example, Mathematical Sciences Publishers is a non-profit, low-cost publisher whose journals have grown in popularity in large part because of the good name of its publisher. Also becoming more common are open access journals, and many standard journals nowadays have an open access option. There is a large ongoing debate about how academic publishing should work that we are barely skimming the surface of here.

With experience, there will usually be a small handful of natural candidate journals to submit to. It is a good strategy to aim for a more prestigious journal first, though don't take it personally if your paper is not accepted. It often takes two or three attempts before your work is accepted. Take any reviewer feedback seriously and try to improve your paper at each stage.

10. Final advice for a researcher

–Some useful resources

- **arxiv.org** Most mathematics papers are initially uploaded here (as *preprints*) and hence publicly available. Once you produce work of your own, it's a good idea to also upload your work here.
- **mathscinet.ams.org** This is the online location of Math Reviews, a service provided by the American Mathematical Society. It is a catalogue of virtually every published math paper, along with a summary of the content and citing/cited references. It provides a very easy way to navigate through the mathematical literature.
- **overleaf.com** LaTeX has become the standard language for writing math papers. Overleaf is an online system for preparing documents in LaTeX and sharing them with collaborators. It also has resources for learning to use LaTeX. (Note that if a project has three or more collaborators, you will need to use the paid version of the service.)

- **mathoverflow.net** This is a popular question/answer forum intended for professional mathematicians. There is a companion site, **math.stackexchange.com**, geared at student-level questions.
- **terrytao.wordpress.com** This is the math blog of Terry Tao, one of the world's foremost mathematicians. In addition to its mathematical content, it is full of career advice and tips for writing math. It also contains links to a large number of other popular math blogs and similar websites.

–Delayed gratification
If you're going to make a career as a researcher, then it will probably be because you have original insights and ideas into problems that no one else does. How does this happen? By thinking deeply about a topic over a long period of time. It has happened multiple times in my still-short research career that I encounter a research problem and immediately recognize how to solve it because of something I had learned years earlier just to satisfy my curiosity.

Make sure you really understand what you're studying and you are not just going through the motions. This may be slower going in the short term but it will pay dividends in the long run. This can be difficult as a graduate student, when layers upon layers of technicalities are getting thrown at you in your coursework or you're attending abstruse seminars in topics you know little about. Do your best. If you feel overwhelmed, take a step back to what you already know and are comfortable with and build upon that.

–Networking
Networking will be essential to your mathematical career. It is hard to get a job offer from a university without people there who know you and are willing to advocate for you. It's also the fun side of mathematics: to travel and socialize with a wide variety of people. This ranges from a simple dinner with an invited seminar speaker to traveling to conferences or university visits of your own.

That said, there is a time and place for everything. Before you can effectively network, first you need to produce some quality research. That way, you'll have something to actually talk about. Otherwise, there is a danger of empty networking, or even reverse networking: to meet many people without making a good impression on them.