

13 Commandments for Writing Mathematics Papers

Writing a mathematics paper is extremely difficult. Writing a clear and readable paper is an order of magnitude more difficult. In these four pages, I will tell you how to do it, by focusing on 13 guiding principles that can be applied to any paper. In case you are looking for explicit examples of these principles, see any of my recent papers; most of them follow most of the principles most of the time.

0. *Time scale.* Decide on how many days you think your paper will take to write. Then multiply that number by 10. The second number is the correct number of days. In order to follow the other principles laid out below, you will have to spend time planning your paper before you start writing. You will need to read and re-read your paper many times, with breaks in between. You will need to completely rewrite sections from scratch. And then there will need to be a cool-down period at the end.

1. *Introductions.* An introduction should tell the reader, as quickly as possible, what is in the paper and why it is important. When you submit your paper to a journal, the editor will ask for quick opinions. Even among these expert reviewers, the rate of readership decays exponentially with the distance from the title. Therefore, do not start your paper with a list of notations or background that the expert reviewer already knows.

Some reasons why your paper is important: it answers a question that was raised elsewhere, it introduces a new construction, it has applications to another sub-area, it raises a new and compelling question, it generalizes another result, it fits into a program, it develops a new theory, or it is surprising. Make sure to explicitly say why your result is interesting; do not expect the reviewers to connect the dots (they won't).

You may find yourself in the situation where the statement of your main result is technical and requires many definitions. In this case, write an introduction to the introduction, that is, an executive summary, and explain the details later (I learned about the intro to the intro from Joan Birman). See any of my recent papers for examples.

One other thing about introductions: explain in very broad terms what are the main new ideas in the paper. You do not want the reader to have to get to Section 6 that you are defining a novel object, or using a familiar object in an unexpected way. Most papers have, at their heart, one key new idea and then a lot of details. Make sure this new idea is brought to light! Your paper will more likely be appreciated, and your work will more likely be used by others.

Oh, and one other thing about introductions: write it in such a way that it gives an overview of and the flavor of the whole paper. In particular, this means describing in broad strokes the arc of the paper and the outline of the proof, plus a description of how the paper is broken up and what is in the various sections.

And one final thing, circling back to the first point about introductions: make sure to state your main result in the introduction, hopefully on the first page or so. If there are too many definitions to state it precisely on the first page, then give the idea in the introduction to the introduction. Remember about the exponential decay. So make sure the reader knows why your paper exists!

2. *Sections*. As a general rule, each section of your paper should contain exactly one main proposition, stated clearly at the beginning, and proved throughout the section, with the help of a series of lemmas. The propositions should piece together in an understandable way in order to prove the theorem. The introduction should explain what is in each of the sections, so that the reader has an overview before they dive into Section 2.

It is helpful to think of your paper as an expandable/hyperlinked document, even though it is not. If you use the method where sections correspond to propositions, then the reader could get a good overview of your paper by reading the introduction to each section. And then if there is a piece they want to dive into, they read that section.

3. *Forward progress*. When we teach Calculus students how to write epsilon-delta proofs, we tell them to find the delta based on the epsilon, and then reverse the steps so that their proof makes forward progress from the beginning to the end. You should do this in your paper, on both the global and local levels. It causes too much mental labor for the reader to have to reverse the argument themselves.

On the global level, the propositions that make up your sections should go forward, not backward. On the local level, the propositions should be proved in the forward direction. And on an even more local level, the individual arguments should go in the forward direction as much as possible. For this reason, avoid proofs by contradiction when possible (it is always possible).

4. *Signposting*. At any given point in the paper, the reader should know why you are proving a certain proposition or lemma. At the beginning of each section, you should remind the reader where you are in the progress of proving your theorem, state the proposition for that section, and remind the reader why this proposition moves the ball

forward. You should then outline the strategy for proving the proposition in the section. If the proof of the proposition requires several moving parts or several steps, then these should be laid out in the beginning. It may be necessary in the middle of a section or a proof to give signposting for where you are in the argument.

5. *Spaghetti code*. According to Wikipedia: Spaghetti code is a pejorative phrase for unstructured and difficult-to-maintain source code. If you take the proof of a theorem and break it down into many lemmas, and then group the lemmas according to topic, you are likely to end up with spaghetti code (see, for example, the typical textbook on abstract algebra). The result is a paper that is hard to read: the forest cannot be seen for the trees. If you follow the previous three principles, you are unlikely to end up with spaghetti code.

6. *No flipping*. This is closely related to the previous principle. If the reader needs to flip back and forth in order to follow the paper (to find a definition, to recall the statement of a lemma, to find where a given proposition fits into the proof of a theorem), then this is a symptom of spaghetti code. A math paper should be able to be read from beginning to end with minimal flipping.

7. *Background sections*. Do not have them. Basically ever. For the expert reader, this is a waste of space (remember the exponential decay). Also, there are only one or two experts on the topic of your paper. For the non-expert reader, you are asking them to take in a lot of material they probably don't know, and carry it around with them as they read your paper. That will make it hard on the reader and unlikely that they will continue reading. Having a background section is an easy way to make your paper self-contained, but unreadable.

8. *Just-in-time references*. If there is no background section, where should you put the background? Answer: just in time. If there is a point where you need to use the fact that the curve graph is hyperbolic, you should write: "Because the curve graph is hyperbolic [MM1], it has thin triangles, and so the theorem is proved." This does the job perfectly. The expert is not slowed down, and the non-expert is informed where to find the proof of the result that is being used.

A corollary is that there is usually no reason to have others' results in a theorem environment (because you will state their result just in time). An added benefit: since most/all of your results in the theorem environment will be yours, this highlights your work better.

9. *Chunking*. According to self-help guru Tony Robbins, chunking is “the grouping together of information into ideally sized pieces so they can be used effectively to produce the outcome you want without stress or shutdown.” This principle applies in spades to mathematics papers. Once you have a proof of a theorem, you should think about what the big pieces are, and then think about what the pieces of the pieces are, etc. There is no right answer, and in general there are many, many ways to chunk.

Proofs of lemmas and propositions should be chunked almost always. Figure out what the main steps of the proof are and break up the proof accordingly. A series of claims can make a proof much easier to digest.

Chunking is an art, not a science. But coming up with the right chunking will greatly illuminate your work for both you and the reader. Sometimes the chunks will have names (like the well-suited curve criterion); this helps to give each chunk a life of its own, and increases the likelihood that others can make use of your ideas.

10. *Scannability*. When I sit down to read a paper, usually what I try to do first is to figure out for myself how the paper is organized. The same goes for the proof of a proposition: I first try to figure out what the overall strategy is (say, proof by cases) and then make a map of the arc of the argument (say, where the cases start and end). In other words, I try to figure out what the chunks are.

You can take care of this step for your reader by making your paper scannable. For instance, if you start the proof by saying it is a proof by cases, and then put the cases in an itemized list, and then put each case in a paragraph with an italicized header, and with whitespace between the paragraphs, the reader can sense what is going on with little effort. Similarly if there are clearly visible claims with whitespace in between.

There are many other ways to make your paper scannable, but your main weapons here are bulleted lists, paragraph labels, and white space. There are other techniques, such as flow charts and diagrams. If you have a page that is a wall of text, think about how you can make it more scannable. A well-written paper can transmit information to a reader that simply flips through, reading section titles, paragraph headers, etc.

11. *Tightening*. Once you have a draft of your paper, make a flow chart of the lemmas, propositions, and theorems. You can decorate your flow chart by circling the results that lie in the same section. Ask yourself: Where is each hypothesis/lemma/definition used? Does it make sense to combine or split any lemmas? You might find that a lemma or

definition is not used. Or you might find that a certain hypothesis is used in only one place, in which case you can try to remove the hypothesis.

12. *Nothing is clear.* Never say in your paper that some fact is easy, clear, obvious, etc. If it really is easy, that usually means that the argument is one sentence; just give the sentence. Sometimes, you can instead just remove the offending phrase; instead of “it is obvious that cyclic groups are abelian,” just say “cyclic groups are abelian.”

If I want to find a mistake in a paper, I scan through for one of these offending phrases. It has happened multiple times that I have found a fatal mistake in a paper this way.

13. Be a scholar. Mainly, what I mean here is doing a thorough literature search to find all closely related papers, and putting your paper in historical context in your introduction. This might be a lot of work, maybe even a week’s worth of work. It is worth the effort, for both your benefit and the reader’s benefit. No theorem is an island - we are all working in the context of the math that came before us. It is important to acknowledge this so that your paper makes sense in context.

While we are on the topic of references, here’s a pro-tip: when you submit your paper to a journal, they will ask for quick opinions. Who are they going to ask? Most of the time it will be people from your introduction or your references. If there is someone that you think would give a good quick opinion, make sure they are mentioned!

Also on this topic, it is in your best interest to be generous to others that came before you. This is for your own emotional health and also so as to not upset others in your field. You will get full credit for your theorem no matter how much credit you give to others, so do not hold back in giving credit wherever you found help or inspiration.

The 13 commandments can be summarized by two meta-commandments:

1. organize your paper well, and
2. make it clear how your paper is organized.

The main ingredient to making these things happen is good old fashioned elbow grease. After you are done with the writing, you should read your paper until you can’t stand it anymore. Force yourself to read the words on the page as if you have never seen them before. Don’t believe any sentence unless you can reproduce the logic on the spot. This is taxing, but possible. This process will greatly improve your paper.

Finally, ask yourself this question: *Would I want to read this paper?* We have all been frustrated trying to read papers. Think about whether you are falling into any of the traps that you notice when reading other papers.

Good luck. I look forward to reading your well-written paper!