


How to Write Mathematics

Paul R. Halmos

L'Enseignement mathématique 16:123–152 (1970)

Contents

1 Preface	1
2 There is no recipe and what it is	2
3 Say Something	2
4 Speak to someone	3
5 Organize first	4
6 Think about the alphabet	6



1 Preface

This is a subjective essay, and its title is misleading; a more honest title might be HOW I WRITE MATHEMATICS. It started with a committee of the American Mathematical Society, on which I served for a brief time, but it quickly became a private project that ran away with me. In an effort to bring it under control I asked a few friends to read it and criticize it. The criticisms were excellent; they were sharp, honest, and constructive; and they were contradictory. “Not enough concrete examples” said one; “don’t agree that more concrete examples are needed” said another. “Too long” said one; “maybe more is needed” said another. there are traditional (and effective) methods of minimizing the tediousness of long proofs, such as breaking them up in a series of lemmas” said one. “One of the things that irritates me greatly is the custom (especially of beginners) to present a proof as a long series of elaborately stated, utterly boring lemmas” said another. There was one thing that most of my advisors agreed on; the writing of such an essay is bound to be a thankless task. Advisor 1: “By the time a mathematician has written his second paper, he is convinced he knows how to write papers, and would react to advice with impatience.” Advisor 2: “All of us, I think, feel secretly that if we but bothered we could be really first rate expositors. People who are quite modest about their mathematics will get their dander up if their ability to write well is questioned.” Advisor 3 used the strongest language; he warned me that since I cannot possibly display great intellectual depth in a discussion of matters

of technique, I should not be surprised at “the scorn you may reap from some of our more supercilious colleagues.” My advisors are established and well known mathematicians. A credit line from me here wouldn’t add a thing to their stature, but my possible misunderstanding, misplacing, and misapplying their advice might cause them annoyance and embarrassment. That is why I decided on the unscholarly procedure of nameless quotations and the expression of nameless thanks. I am not the less grateful for that, and not the less eager to acknowledge that without their help this essay would have been worse. ”Hier stehe ich; ich kann nicht anders.”

2 There is no recipe and what it is

I think I can tell someone how to write, but I can’t think who would want to listen. The ability to communicate effectively, the power to be intelligible, is congenital, I believe, or, in any event, it is so early acquired that by the time someone reads my wisdom on the subject he is likely to be invariant under it. To understand a syllogism is not something you can learn; you are either born with the ability or you are not. In the same way, effective exposition is not a teachable art; some can do it and some cannot. There is no usable recipe for good writing. Then why go on? A small reason is the hope that what I said isn’t quite right; and, anyway, I’d like a chance to try to do what perhaps cannot be done. A more practical reason is that in the other arts that require innate talent, even the gifted ones who are born with it are not usually born with full knowledge of all the tricks of the trade. A few essays such as this may serve to “remind” (in the sense of Plato) the ones who want to be and are destined to be the expositors of the future of the techniques found useful by the expositors of the past. The basic problem in writing mathematics is the same as in writing biology, writing a novel, or writing directions for assembling a harpsichord: the problem is to communicate an idea. To do so, and to do it clearly, you must have something to say, and you must have someone to say it to, you must organize what you want to say, and you must arrange it in the order you said it in, you must write it, rewrite it, and re-rewrite it several times, and you must be willing to think hard about and work hard on mechanical details such as diction, notation, and punctuation. That’s all there is to it.

3 Say Something

It might seem unnecessary to insist that in order to say something well you must have something to say, but it’s no joke. Much bad writing, mathematical and otherwise, is caused by a violation of that first principle. Just as there are two ways for a sequence not to have a limit (no cluster points or too many), there are two ways for a piece of writing not to have a subject (no ideas or too many). The first disease is the harder one to catch. It is hard to write many words about nothing, especially in mathematics, but it can be done, and the result is bound to be hard to read.

There is a classic crank book by Carl Theodore Heisel [5] that serves as an example. It is full of correctly spelled words strung together in grammatical sentences, but after three decades of looking at it every now and then I still cannot read two consecutive pages and make a one-paragraph abstract of what they say; the reason is, I think, that they don’t say anything. The second disease is very common: there are many books that violate the principle of having something to say by trying to

say too many things. Teachers of elementary mathematics in the U.S.A. frequently complain that all calculus books are bad. That is a case in point. Calculus books are bad because there is no such subject as calculus; it is not a subject because it is many subjects. What we call calculus nowadays is the union of a dab of logic and set theory, some axiomatic theory of complete ordered fields, analytic geometry and topology, the latter in both the “general” sense (limits and continuous functions) and the algebraic sense (orientation), real-variable theory properly so called (differentiation), the combinatoric symbol manipulation called formal integration, the first steps of low-dimensional measure theory, some differential geometry, the first steps of the classical analysis of the trigonometric, exponential, and logarithmic functions, and, depending on the space available and the personal inclinations of the author, some cook-book differential equations, elementary mechanics, and a small assortment of applied mathematics. Any one of these is hard to write a good book on; the mixture is impossible. Nelson’s little gem of a proof that a bounded harmonic function is a constant [7] and Dunford and Schwartz’s monumental treatise on functional analysis [3] are examples of mathematical writings that have something to say. Nelson’s work is not quite half a page and Dunford-Schwartz is more than four thousand times as long, but it is plain in each case that the authors had an unambiguous idea of what they wanted to say. The subject is clearly delineated; it is a subject; it hangs together; it is something to say. To have something to say is by far the most important ingredient of good exposition—so much so that if the idea is important enough, the work has a chance to be immortal even if it is confusingly misorganized and awkwardly expressed. Birkhoff’s proof of the ergodic theorem [1] is almost maximally confusing, and Vanzetti’s “last letter” [9] is halting and awkward, but surely anyone who reads them is glad that they were written. To get by one the first principle alone is, however, only rarely possible and never desirable.

4 Speak to someone

The second principle of good writing is to write for someone. When you decide to write something, ask yourself who it is that you want to reach. Are you writing a diary note to be read by yourself only, a letter to a friend, a research announcement for specialists, or a textbook for undergraduates? The problems are much the same in any case; what varies is the amount of motivation you need to put in, the extent of informality you may allow yourself, the fussiness of the detail that is necessary, and the number of times things have to be repeated. All writing is influenced by the audience, but, given the audience, an author’s problem is to communicate with it as best he can. Publishers know that 25 years is a respectable old age for most mathematical books; for research papers five years (at a guess) is the average of obsolescence. (Of course there can be 50-year old papers that remain alive and books that die in five.) Mathematical writing is ephemeral, to be sure, but if you want to reach your audience now, you must write as if for the ages. I like to specify my audience not only in some vague, large sense (e.g., professional topologists, or second year graduate students), but also in a very specific, personal sense. It helps to think of a person, perhaps someone I discussed the subject with two years ago, or perhaps a deliberately obtuse, friendly colleague, and then to keep him in mind as I write. In this essay, for instance, I am hoping to reach mathematics students who are near the beginning of their thesis work, but, at the same time, I am keeping my mental eye on a colleague whose ways can stand mending. Of course I hope that (a)

he'll be converted to my ways, but (b) he won't take offense if and when he realizes that I am writing for him. There are advantages and disadvantages to addressing a very sharply specified audience. A great advantage is that it makes easier the mind reading that is necessary; a disadvantage is that it becomes tempting to indulge in snide polemic comments and heavy-handed "in" jokes. It is surely obvious what I mean by the disadvantage, and it is obviously bad; avoid it. The advantage deserves further emphasis. The author must anticipate and avoid the reader's difficulties. As he writes, he must keep trying to imagine what in the words being written may tend to mislead the read, and what will set him right. I'll give examples of one or two things of this kind later; for now I emphasize that keeping a specific reader in mind is not only helpful in this aspect of the writer's work, it is essential. Perhaps it needn't be said, but it won't hurt to say, that the audience actually reached may differ greatly from the intended one. There is nothing that guarantees that a writer's aim is always perfect. I still say it's better to have a definite aim and hit something else, than to have an aim that is too inclusive or too vaguely specified and no chance of hitting anything. Get ready, aim, and fire, and hope that you'll hit a target: the target you were aiming at, for choice, but some target in preference to none.

5 Organize first

The main contribution that an expository writer can make is to organize and arrange the material so as to minimize the resistance and maximize the insight of the reader and keep him on the track with no unintended distractions. What, after all, are the advantages of a book over a stack of reprints? Answer: efficient and pleasant arrangement, emphasis where emphasis is needed, the indication of interconnections, and the description of the examples and counterexamples on which the theory is based; in one word, organization. The discoverer of an idea, who may of course be the same as its expositor, stumbled on it helter-skelter, inefficiently, almost at random. If there were no way to trim, to consolidate, and to rearrange the discovery, every student would have to recapitulate it, there would be no advantage to be gained from standing "on the shoulders of giants", and there would never be time to learn something new that the previous generation did not know. Once you know what you want to say, and to whom you want to say it, the next step is to make an outline. In my experience that is usually impossible. The ideal is to make an outline in which every preliminary heuristic discussion, every lemma, every theorem, every corollary, every remark, and every proof are mentioned, and in which all these pieces occur in an order that is both logically correct and psychologically digestible. In the ideal organization there is a place for everything and everything is in its place. The reader's attention is held because he was told early what to expect, and, at the same time and in apparent contradiction, pleasant surprises keep happening that could not have been predicted from the bare bones of the definitions. The parts fit, and they fit snugly. The lemmas are there when they are needed, and the interconnections of the theorems are visible; and the outline tells you where all this belongs. I make a small distinction, perhaps an unnecessary one, between organization and arrangement. To organize a subject means to decide what the main headings and subheadings are, what goes under each, and what are the connections among them. A diagram of the organization is a graph, very likely a tree, but almost certainly not a chain. There are many ways to organize most subjects, and usually there are many ways to arrange the results of each method of organization in a linear order. The organization is more

important than the arrangement, but the latter frequently has psychological value. One of the most appreciated compliments I paid an author came from a fiasco; I botched a course of lectures based on his book. The way it started was that there was a section of the book that I didn't like, and I skipped it. Three sections later I needed a small fragment from the end of the omitted section, but it was easy to give a different proof. The same sort of thing happened a couple of times more, but each time a little ingenuity and an ad hoc concept or two patched the leak. In the next chapter, however, something else arose in which what was needed was not a part of the omitted section but the fact that the results of that section were applicable to two apparently very different situations. That was almost impossible to patch up, and after that chaos rapidly set in. The organization of the book was tight; things were there because they were needed; the presentation had the kind of coherence which makes for ease in reading and understanding. At the same time the wires that were holding it all together were not obtrusive; they became visible only as part of the structure was tampered with. Even the least organized authors make a coarse and perhaps unwritten outline; the subject itself is, after all, a one-concept outline of the book. If you know that you are writing about measure theory, then you have a two-word outline, and that's something. A tentative chapter outline is something better. It might go like this: I'll tell them about sets, and then measures, and then functions, and then integrals. At this stage you'll want to make some decisions, which, however, may have to be rescinded later; you may for instance decide to leave probability out, but Haar measure in. There is a sense in which the preparation of an outline can take years, or, at the very least, many weeks. For me there is usually a long time between the first joyful moment when I conceive the idea of writing a book and the first painful moment when I sit down and begin to do so. In the interim, while I continue my daily bread and butter work, I daydream about the new project, and, as ideas occur to me about it, I jot them down on loose slips of paper and put them helter-skelter in a folder. An "idea" in this sense may be a field of mathematics I feel should be included, or it may be an item of notation; it may be a proof, it may be an aptly descriptive word, or it may be a witticism that, I hope, will not fall flat but will enliven, emphasize, and exemplify what I want to say. When the painful moment finally arrives, I have the folder at least; playing solitaire with slips of paper can be a big help in preparing the outline. In the organization of a piece of writing, the question of what to put in is hardly more important than what to leave out; too much detail can be as discouraging as none. The last dotting of the last i, in the manner of the old-fashioned Cours d'Analyse in general and Bourbaki in particular, gives satisfaction to the author who understands it anyway and to the helplessly weak student who never will; for most serious-minded readers it is worse than useless. The heart of mathematics consists of concrete examples and concrete problems. Big general theories are usually afterthoughts based on small but profound insights; the insights themselves come from concrete special cases. The moral is that it's best to organize your work around the central, crucial examples and counterexamples. The observation that a proof proves something a little more general than it was invented for can frequently be left to the reader. Where the reader needs experienced guidance is in the discovery of the things the proof does not prove; what are the appropriate counterexamples and where do we go from here?

6 Think about the alphabet

Once you have some kind of plan of organization, an outline, which may not be a fine one but is the best you can do, you are almost ready to start writing. The only other thing I would recommend that you do first is to invest an hour or two of thought in the alphabet; you'll find it saves many headaches later. The letters that are used to denote the concepts you'll discuss are worthy of thought and careful design. A good, consistent notation can be a tremendous help, and I urge (to the writers of articles too, but especially to the writers of books) that it be designed at the beginning. I make huge tables with many alphabets, with many fonts, for both upper and lower case, and I try to anticipate all the spaces, groups, vectors, functions, points, surfaces, measures, and whatever that will sooner or later need to be baptized. Bad notation can make good exposition bad and bad exposition worse; ad hoc decisions about notation, made mid-sentence in the heat of composition, are almost certain to result in bad notation. Good notation has a kind of alphabetical harmony and avoids dissonance. Example: either $+by$ or $+a_2x_2$ is preferable to $+bx_2$. Or: if you must use Σ for an index set, make sure you don't run into $\sum_{\sigma \in \Sigma} a_\sigma$. Along the same lines: perhaps most readers wouldn't notice that you used $|z|$.