

How to talk Mathematics

P.R. Halmos

Apology

The purpose of what follows is to suggest to a young mathematician what he might do (and what he had better not do) the first few times that he gives a public lecture on his subject. By a "public lecture" I mean something like a colloquium talk (to more or less the entire mathematics department at a large university), or an invited address (to more or less the entire membership on attendance at a meeting of the American Mathematical Society); I do not mean a classroom lecture (to reluctant beginners) or a seminar talk (to dedicated experts). That an article on how to talk mathematics might serve a good purpose was suggested by the officers of the American Mathematical Society. It seems that there have been more and more complaints about invited addresses ("they are incomprehensible, and therefore useless"), and that, therefore, it might do some good to let a speaker know about such complaints before he adds to the reasons for them.

A genius makes his own rules, but a "how to" article is written by one ordinary mortal for the benefit of another. Harpo Marx, one of the greatest harpist of all times, was never taught how to play; everything he did was "wrong" accordingly to standard teaching. Most things that an article such as this one can say have at least a counterexample in the practice of some natural born genius. Authors of articles such as this one know that, but in the first approximation they must ignore it, or nothing would ever get done.

Why lecture?

What is the purpose of a public lecture? Answer: to attract and to inform. We like what we do, and we would like for others to like it too; and we believe that the subject's intrinsic qualities are good enough so that anyone what they are cannot help being attracted to them. Hence, better answer: the purpose of a public lecture is to inform, but to do so in a manner that makes it possible to the audience to absorb the information. An attractive presentation with no contents is worthless, to be sure, but a lump of indigestible information is worth no more.

The question then becomes this: what is the best way to describe a subject (or that small part of a subject that has recently been the center of the lecturer's attention) to an audience of mathematicians most of whom are interested in something else? The problem is different from describing a subject to students who must learn it in usable detail, and it is different from sharing a new discovery with fellow experts who have been thinking about the same sort of things and are wondering what you know what they don't.

Simplicity

Less is more, said the great architect Mies van der Rone, and if all lecturers remembered that adage, all audiences would be both wiser and happier.

Have you ever disliked a lecture because it was too elementary? I am sure that there are people who would answer yes to that question, but not many. Every time I have asked the question the person who answered said no, and then looked a little surprised at hearing the answer. A public lecture should be simple and elementary; it should not be complicated and technical. If you believe and can act on this injunction ("be simple") you can stop reading here; the rest of what I have to say is, in comparison, just a matter of mirror detail.

To begin with a public lecture to 500 people with "Consider a sheaf of germs of holomorphic functions ..." (I have heard it happen) loses people and antagonizes them. If you mention the Kimneth formula, it does no harm to say that, at least as far as Betti numbers go, it is just like what happens when you multiply polynomials. If you mention functors, say that a typical example is the formation of the duals of vector spaces and the adjoints of linear transforms. Be simple by being concrete. Listeners are prepared to accept unstated (but hinted) generalizations much more than they are able, on the spur of the moment, to decode a precisely stated abstraction and to re-invent the special cases that motivated it in the first place. Caution: being concrete should not lead to concentrate seeing the trees and not seeing the woods. In many parts of mathematics a generalization is simpler and more incisive than its special parent. (Examples: Artin's solution of Hilbert's

17th problem about definite forms via formally real fields; Gelfand's proof of Wiener's theorem about absolutely convergent Fourier series via Banach algebras). In such cases is always a concrete special case that is simpler than the seminal one and that illustrates the generalization with less fuss; the lecturer who knows his subject will explain the complicated special case, and the generalization, by discussing the simple cousin.

Some lecturers defend complications and technicalities by saying that that's what their subject is like, and there is nothing they can do about it. I am skeptical, and I am willing to go so far as to say that such statements indicate incomplete understanding of the subject and of its place in mathematics. Every subject and even every small part of a subject, if it is identifiable, if it is big enough to give an hour talk on, has its simple aspects, and they, the simple aspects, the roots of the subject, the connections with more widely known and older parts of mathematics are what a non-specialized audience needs to be told.

Many lecturers, especially those near of the foot of the academic ladder, anxious to climb rapidly, feel under pressure to say something brand-new - to impress their elders with their brilliance and profundity. Two comments: (1) the best way to do that is to talk simple, and (2) it doesn't really have to be done. It may be entirely appropriate to make the lecturer's recent research the focal point of the lecture, but it may also be entirely appropriate not to do so. An audience's evaluation of the merits of a talk is not proportional to the amount of original material included; the explanation of the speaker's latest theorem may fail to improve his chances of creating a good impression.

An oh-quoted compromise between trying to be intelligible and trying to seem deep is this advice: address the first quarter of your talk to your high-school chemistry teacher, the second to a graduate student, the third to an educated mathematician whose interests are different from yours, and the last to the specialists. I have done my duty by reporting the formula, but I'd fail in my duty if I didn't warn that there are many who do not agree with it. A good public lecture should be a work of art. It should be an architectural unit whose parts reinforce each other in conveying the maximum possible amount of information -not a campaign speech that offers something to everybody and, more like than not, ends by pleasing nobody.

Make it simple, and you won't go wrong.

Details

Some lecturers, with the best of intentions, striving for simplicity, try to achieve it by being overly explicit and overly detailed; that's a mistake.

"Explicit" refers to computations. If a proof can be carried by multiplying two horrendous expressions, say it so and let it go at that; the logical simplicity of the steps doesn't necessarily make the computations attractive or informative to carry out. Landau, legend has it, never omitted a single epsilon of his lectures, and his lectures were inspiring anyway -but that's the exception, not the rule. If, on an exceptional occasion, you decide that a brief computation will be decisive and illuminating, put it in, but the rule for ordinary mortals still stands: do not compute in public. It may be an explicit and honest thing to do, but that's not what makes a lecture simple.

"Detailed" refers to definitions. Some lecturers think that the way to reach an audience of non-experts is to tell them everything. ("To get to the theorem I proved last week, I need starting from the beginning, 14 definitions and 11 theorems that my predecessors have proved. If I talk and write fast, I can present those 25 nuggets in 25 minutes, and in the rest of the time I can state and present my own thing.") This, too, is honest, and it makes the lecture self-contained, in some sense -but it is impossible to digest, and the effect is dreadful. If someone told you, in half an hour, the meaning of each ideogram on a page of Chinese, could you then read and enjoy the poem on that page in the next half hour?

Proofs

Some lecturers understand the injunction "be simple" to mean "don't prove anything". That isn't quite right. It is true, I think, that it is not the main purpose of a public lecture to prove things, but to prove nothing robs the exposition of an essential part of what mathematicians regard as attractive and informative. I would advise to any lecturer to be sure to prove something -one little theorem, one usable and elegant lemma, something that is typical of the words and the methods used in the subject. If the proof is short enough, it almost doesn't matter that it may, perhaps, not be understood. It is of value to the listener to hear the lecturer say that Bernoulli numbers enter the theory of stable homotopy groups,

even if the listener has only an approximate idea of what Bernoulli numbers or homotopy groups are.

Something that's even better than a sample proof is the idea of the proof, the intuition that suggests it in the first place, the reason why the theorem is true. To find the right words to describe the central idea of a proof is something hard, but it is worth the trouble; when it can be done, it provides the perfect way to communicate mathematics.

Problems

In the same vein, it is a false concept of simplicity that makes a lecturer concentrate only in what is safe and known; I strongly recommend that every public lecture reach the frontiers of knowledge, and at least mention something that is challenging and unknown. It doesn't have to be, it shouldn't be, the most delicate and newest technicality. Don't be afraid of repeating an old one; remember that many of your audience haven't heard of your subject since they took a course in it in graduate school, a long time ago. They will learn something just by hearing today that the unsolved problem they learned about years is still unsolved. The discussion of unsolved problems is a valuable part of the process of attracting and informing -it is, I think, an indispensable part. A field is not well described if its boundaries are missing from the description; some knowledge of the boundaries is essential for the understanding of where the field is today as well as for enlarging the area of our knowledge tomorrow. A public lecture must be simple, yes, but not at the cost of being empty, or, not quite that bad but bad enough; it must not be incomplete to the point of being dishonest.

Organization

The organization of a talk is like the skeleton of a man: things would fall apart without it, but it's bad if it shows. Organize your public lecture, plan it, prepare it carefully, and then deliver it impromptu, extemporaneously.

To prepare a talk, the first thing to know is the subject, and very close second is the audience. It's much more important to adjust the level to fit the audience in a public lecture than it is in a book. ("Adjust the level" is not an euphemism for "talk down". Don't insult the audience, but be realistic. Slightly over the mark, very slightly, doesn't do much harm, but too much over is much worse than something under). A reader can put down a book, and can come back to it when he has learned more; an antagonized listener will, in spirit, leave you, and, as far as this talk is concerned, he'll never come back. The right level for a talk is a part of what organization is meant to achieve, but, of course, the first and more important thing to organize is the content. Here I have two recommendations (in addition to "prove something" and "ask something", already mentioned): (1) discuss three or four related topics, and the connections between them, rather than relentlessly pursue one central topic, and (2) break each topic in four or five sub-topics, portable, freely addable or subtractable modules, the omission of any of which would not wreck the continuity. As for extemporaneous delivery, there are two reasons for that: it sounds good, and it makes possible an interaction between the speaker and the listeners. The faces in the audience can be revealing and helpful: they can indicate the need to slow down, to speed up, to explain something, to omit something.

Preparation

To prepare a lecture means to prepare the subjects it will cover, the order in which those subjects are to come, and the connections between them that you deem worthy to mention; it does not mean to write down all the words with the intention of memorizing them (or, much worse, reading them aloud). Still: to write it all out is not necessarily a bad idea. "All" means all, including, especially, exactly what is to be put on the blackboard (with a clear idea of when it will be put on and whether it will remain for a long or be rubbed out right away). To have it all written out will make it easier to run through it once, out loud, by a blackboard, and thus to get an idea of the timing. (Warning: if the dry run takes an hour, then the actual delivery will take an hour and a half.)

Brevity

Most talks are described as "one-hour lectures", but, by a generally shared tradition, most are meant to last for 50 minutes only. Nobody will reproach you for sitting down after 45 minutes, but the majority of the audience will become nervous after 55, and most of them will glare at you, displeased and uncomfortable, after 65. To take long, to run over time, is rude. Your theorems, or your proofs, are not all that important in other people's lives; that hurried, breathless, last five minutes is expendable. If you didn't finish, say so, express your regret if you must, but stop; it's better thus than to give the audience cause for regret.

Techniques

A public lecture usually begins with an introduction by the chairman of the session. Rule of etiquette: give him a chance. Before the lecture begins, sit somewhere by the side of the room, or with the audience, near the front; do not stand by or near the blackboard, or hover near the chairman worrying him.

One good trick to overcome initial stage-fright is to memorize one sentence, the opener. After that, the preparation and your knowledge of the subject will take over. Try very hard to avoid annoying mannerisms. Definition: an annoying mannerism is anything that's repeated more than twice. A mannerism can be verbal ("in other words" pronounced " 'n 'zer w'rs", meaning nothing); it can be visual (surrounding a part of the material on the blackboard by elaborate fences); or it can be dynamic (teeter-tottering at the edge of the platform).

If you are in mechanical trouble, catch the chairman's eye and say, to him only, "I am out of chalk", or "May I have an eraser?". Do not bumble about your awkwardness and do not keep on apologizing. ("Oh, dear, where can I put this -sorry, I seem to have run out of the room - well, let's see perhaps we don't need this anymore ♦ ".) Make the appropriate decision and take the appropriate action, but do so silently. Keep your own counsel, and do not distract the audience with irrelevancies.

Silence is a powerful tool at other times too; the best speakers are also the best non-speakers. A long period of silence (five seconds, say, or ten at most) after an important and crisply stated definition or theorem puts the audience on notice ("this is important") and gives them a chance to absorb what was just said. Don't overdo it, but three or four times during the hour, at the three or four high points, you might very well find that the best way to explain something is to say nothing.

Speak slowly and speak loudly; write large and speak as you write; write slowly and do not write much. Intelligently chosen abbreviations, arrows for implications, and just remainder words, not deathless prose, are what a board is for; their purpose is to aid the audience in following you by giving them something to look at as well as something to listen to. (Example: do not write "semisimple is defined as follows "; write "semisimple: ".) Do not, ever, greet an audience with a carefully prepared blackboard (or overhead projector sheets) crammed with formulas, definitions, and theorems. (An occasionally advisable exception to this rule have to do with pictures -if a picture, or two pictures, would help your exposition but would take too long to draw as you talk, at least with the care it deserves, the audience will forgive you for drawing it before the talk begins.) The audience can take pleasure in seeing the visual presentation grow before its eyes - the growth is part of your lecture, or should be.

Flexibility

Because of the unpredictability of the precise timing (you didn't rehearse enough, the audience asks questions during the talk, the lecture's room is reserved for another lecture at 5:00 sharp, or you just plain get mixed and waste time trying to get unscrambled), flexibility is an important quality to build into a lecture. You must be prepared to omit (or to add) material, and you must be prepared to do so under pressure, in public, on the spur of the moment, without saying so, and without seeming to do so. There are probably many ways to make a lecture flexible. I'll mention two that I have found useful. The first is exercises. Prepare two or three statements whose detailed discussion might well be a part of the lecture, and whose omission would not destroy continuity, and, at the proper place during the your lecture, "assign" them to the audience as exercises. You run the risk of losing the attention of some of the more competitive members of the group♦ for the rest of the hour. What you gain is something else that you can gracefully fill out your time with it (unlikely as that may be, you finish everything else too soon, and, at the same time, something that'll never be missed if you do not discuss the solution. Exercises in this sense may yield another fringe benefit: they'll give the audience something to ask their courtesy questions about).

A second way to make a lecture flexible is, and I mentioned before and I believe is worth emphasizing again: portable modules. My notes of a lecture usually consists of about 20 telegraphically written paragraphs. The detailed presentation of each paragraph may take between 2 and 4 minutes, and at least half the paragraphs (the last 10) are omissible. These omissible modules often contain material dear to my heart: that clever proof, that ingenious generalization, that challenging question -but no one (except me) will miss them if I keep them. Knowing that those modules are there, I sail through the first half of the period with no worries: I am sure that I won't run out of things to say, and I am sure that everything that I

must say will get said. In the second half of the period I have an eye on the time, and, without saying anything about it, make instantaneous decisions about what to throw overboard. One disadvantage of this method is that at the end of your time you may sound too abrupt, as if you have stopped in the middle of a sentence. To avoid the abrupt ending, prepare your peroration, and do not omit it. The peroration can be a three-sentence summary of the whole lecture, or it can be the statement of the most important unsolved problems of the subject. Make whatever you think proper for an ending, and then end with it. Rule of etiquette: when you stop, sit down. Literally sit down. Do not just stop talking and look helpless, and do not ask for questions: that's the chairman's job.

Short talks

Short talks are harder to prepare and to deliver than long ones. The lecturer has less time to lay the groundwork, and the audience is under pressure to understand quickly.

In my experience a 20-minute talk can still be enjoyable and enlightening; all you need is prepare a 10-minute talk and present it leisurely. A 10-minute talk is the hardest to do right; the precepts presented above (simple, organized and short) must be applied again, but this time there is no room for error. Focus on one idea only, and on its simplest non-trivial special case at that, practice the talk and time it carefully, and under no circumstances allow a 10-minute contributed paper to become a 45-minute uninvited address. It has been done, but the results were neither informative nor attractive. Some experts are willing to relax the rules for a 10-minute talk: it is all right, they say, to dive into the middle of things immediately. and it is all right, they say, to use prepared projection sheets. Others, having in mind the limited velocity and capacity of the human mind to absorb technicalities, disagree.

Summary

My recommendations amount to this: make it simple, organized and short. Make your lecture simple (special and concrete); be sure to prove something and ask something; prepare, in detail; organize the content and adjust to the level of the audience; keep it short, and, to be sure of doing so, prepare it so as to make it flexible. Remember that you are talking in order to attract the listeners to your subject and to inform them about it; and remember that less is more.