

THREE RULES FOR LEADING A PAPER DISCUSSION

John M. Drake

February 26, 2013

Reading and discussing research papers is one of the key occasions of learning in college or graduate school. Whether as part of a course, journal club, seminar, lab meeting or study group, the "paper discussion" is at the center of almost any education. Indeed, good paper discussions are one of the key catalysts for scientific innovation and the source of numerous inspired collaborations and breakthroughs. However, some paper discussions are unstimulating, even a waste of time. This is unfortunate because vigorous discussions of research with colleagues can be one of the most enjoyable moments in a life of research. Bad paper discussions are a significant missed opportunity.

What is it that distinguishes the good paper discussions from those that are sterile, forced, or misguided? In my view, the difference is *leadership*. That is, sometimes students (and perhaps professors, too) assume that a discussion will take care of itself, as long as it is given a good start (though, evidently, many think that not even this is needed!). This is wrong. Research articles, like any other piece of communication, are open texts. There is an infinity of ways to engage such a document. There are some crucial first steps: (1) finding out what an author has to say, (2) determining why she would want to say it (to advance a new theory? to test an existing one?), and, particularly important in the interpretation of science, (3) evaluating the evidence presented. In addition, and in science just as in humanistic endeavors, there are a wealth of contextual and critical activities that should be undertaken to "read" a scientific paper. These include digging up its predecessors, envisioning descendant research projects, criticizing methods, entertaining alternate ways the study could have been performed, extending or restricting the scope, enumerating and considering alternative explanations, imagining technological applications, and so forth.

Why are these reflections relevant to leading a paper discussion? It is because such extensive possibilities of reading are unwieldy, especially in the charge of a group of discussants with disparate aims and interests, perhaps some of whom merely skimmed the paper in the first case. What the discussion needs, then, is direction, which can only be provided by a particularly careful reader, a leader, to take charge, to set bounds on the scope and to set the terms of the discussion, to provide logical flow from one idea or concept to the next, to spark conversation when the discussion falters, to enforce standards of aptness and fairness, and (above all) to ensure that the discussion is intellectually productive, that is, to ensure that learning occurs.

These reflections are summarized in a slogan for academic leadership that I take to be axiomatic: the discussion leader's purpose and goal is to ensure that the discussion is an *occasion of learning*. While some discussion leaders are apparently able to pull off such leadership off-the-cuff, for most of us to lead a good discussion requires considerable organization. What follows are some suggestions, based on my observations and experiences, about how to lead a paper discussion effectively.

Rule one: A good discussion leader will start with a summary

First, the commonest mistake when leading a paper discussion is not a misstep, but the failure to take any step at all. A good discussion has a direction, a springboard from which to launch a round of debate, or a foil to criticize. You can bet it will be a bad discussion when the leader's first (typically *only*) remark is, "So, what did you think?" Thus, a good discussion leader will start with a summary.

The summary can be brief. But the key is that it is declarative, not interrogative. It is pointed. It is selective. A summary establishes the leader's control. He or she uses the summary to steer the discussion, but establishing a vocabulary, the working definitions of that vocabulary, and bounds on the aims of the investigation. A summary enables the leader to bring into the discussion those who have not read or have not read carefully (let us call these the "skimmers"). This inclusiveness is important if learning is to be maximized. If the skimmers are not brought in, there will be fewer participants once the group has gotten into the thick of it. Fewer participants is fewer

perspectives, less thinking power, more opportunity for runaway opinion. On the other hand, if the skimmers are brought into the discussion solely or primarily through the leader's summary, they will be allies to the leader later on. (Whether agreeing or disagreeing with the leader, any discussant who takes an interesting and intelligible position is an ally from the standpoint of the leader--whose main goal, we recall, is not so much to find the truth of the matter, but to make the discussion an occasion of learning). In contrast, if the skimmers are not brought in they must, of necessity, be decidedly neutral--and that is a bad thing for a paper discussion. The summary is the opening gambit and it is one of the savvy leader's chief strategic devices.

Rule two: A good discussion leader will plan for the the life cycle of the discussion

At the start, I suggested that a good paper discussion may be a catalyst for scientific innovation. How does this occur? The answer is that the collective understanding that emerges in group discussion may be more complete and less prejudiced than the analysis of any of its participants thinking alone. A good discussion, then, will be one in which a variety of ideas are raised, considered, refined, and organized into a whole analysis, the process of which has been an occasion of learning and the product of which is a critical analysis.

If the first mistake has been avoided, then the discussion will be off to a good start, prompted by the summary. However, after the terms of the discussion have been set and the skimmers have been brought on board, it remains still to bring to the surface the thoughts of the other, non-leader discussants, and, thereafter, to use these to prompt new reflections, which themselves may catalyze further ideas in an avalanche of thinking, which (occasionally) marks the origin of scientific breakthrough. Note that regardless of whether this is a private breakthrough in understanding or the very public sort of breakthrough that leads to an advance in science, in either case the breakthrough is a special moment in a life of research. So, how can such chain reactions of thinking be promoted? The answer is by following the second rule for effectively leading a good paper discussion. A good discussion leader will plan for the the life cycle of the discussion.

Life cycle planning, in this case, involves two key phases: (1) a mechanism to get the chain reaction off the ground, to move from the summary to the autonomous generation of ideas within the group, and (2) a strategy for regulating the discussion as it proceeds, including the ability to provide a jumpstart (should it become necessary) and a framework for wrapping up, an exit strategy.

The mechanism to move from summary to chain reaction is to have prepared good questions. What is a good question? (That is a good question!) In my view, a good question has two properties. First, it has a determinate answer. It is not open-ended. Its answer is verifiable. Understanding of the answer is a test of comprehension. Second, this determinate answer is not too easy. It is not obvious. It cannot be guessed. Naturally, good questions are hard to devise--but they are the life of a discussion! In preparation, the good leader does not squander the opportunity to develop some good questions. Indeed, questions are very useful for all kinds of discussion-leading tasks. Not only are questions useful to spark a new chapter in the discussion, but they can also be used to end one, cutting short the discussion-hogger (who speaks to be heard) and the topic-hijacker (for whom all roads lead to Rome) in a non-confrontational way.

The strategy for regulating discussion is to provide a device. The device may be a rubric (a cross-classified table), a diagram, or a bullet list. It sets the scope of discussion by demonstration rather than by declaration, and gives bones to the meat of the conversation. It is a reference point by which to evaluate if a comment is off topic, a standard to define the terms of discussion, an anchor to keep the discussion in place, and cement to hold it altogether. The good leader never begins a discussion without a device.

Rule three: A good discussion leader will ensure that disagreement is well founded

The final mistake, after failing to summarize and failing to plan, is to confuse criticism with disagreement. It is a beginner's mistake, because unexamined disagreement is intellectually foolhardy and easily exposed by a more careful, more thoughtful discussant. It is silly to question statistical results prior to a considering a study's design

and obtuse to object to a model's assumptions before seeing to what end the model is put. A fair discussant will point this out. Often enough, disagreement from the outset is an excuse for not engaging, for intellectual laziness. If a paper/position/finding can be written off from the start then there's no reason to do the hard thinking-work of analysis and the discussants should be free to leave their minds checked at the door. A good leader does not allow this. A good leader, perhaps gently (perhaps not), does not tolerate bullshit, posturing, or their typical manifestation in conversation--impressive vagueness. That is, a good discussion leader will ensure that disagreement is well founded.

A corollary mistake is to assume that to be critical requires one to be negative. This is worth underscoring: criticism is not the same thing as negativity. I recommend to my students the following three steps for reading a research article. They are also a good guide for a discussion. The steps take the form of questions. The careful reader does not proceed from one step to the next until she can verbalize an answer to the question. The reader should not be deceived by the simplicity of the questions. Simple questions often have difficult answers. Articulating these answers may require some level of subtlety in word choice and logic--in short, critical analysis.

(1) What have the authors found? Notice this is a question about what is reported. The question should be answered with reference to measurements and observations, with as little reference to theoretical terms as possible, especially when the study itself is aimed at testing a theory. It gets answered with statements like the following:

Subjects in treatment group A were X times more likely to present with Y than subjects in treatment group B

or

An increase of X units in variable A corresponded to an increase in Y units of variable B

Of course, these examples are simple. Most scientific findings are not so easily reported. But the goal is the same, no matter how sophisticated the analysis: to transport your mind to the position of the original investigators looking at the data. What did they see? What did they measure?

Of course, seeing and measuring only take one so far in science. Thus, question (2), which is where we start getting in deep with respect to *interpretation*.

(2) How did the authors interpret their findings? This question is crucial for good criticism. Too often, the discussing student wants to jump from findings to her own interpretations. This would be premature, and for two reasons.

First, the authors of the paper have doubtless thought about the meaning of their results far more than anybody else (certainly far more than the skimmer). They are therefore probably pretty good guides to extracting meaning from data, or at least to viewing data in a sophisticated context and in an intellectually responsible way. Yes, the authors may have egos and agendas that cloud their consideration. Guided by their own research interests, intellectual rearing, and pet hypotheses, they may have provincial, outdated, or positively wacky interpretations of their results. But these are uncommon in science. Much more typically, the author's own reflections are more incisive, more penetrating, and more expansive than all but the most prescient readers can muster on a first reading. What's more, the paper has been through peer review, often multiple rounds, and editorial oversight. A reader who places himself or herself in judgement over authors, reviewers, and editors is a bold reader indeed. Authors are not omniscient, but they are good guides to interpretation and should be followed until a careful case has been brought to the contrary--and that doesn't happen until question (3) is asked.

The second reason why the question must be answered is to anticipate what comes in the later stages of critical reading. In the end, readers will often want to accept part of the authors' explanation, or to accept it under different circumstances or with different scope of application, or under the moderation of additional qualifications. All of these are differences of opinion between the authors and reader. To be responsible in holding a different opinion, it seems to me, there must be articulable *reasons* for disagreement. A prerequisite for reasons to disagree is an account of differences, and a prerequisite for that is to understand the contrary opinion. Since the authors wrote their piece first, in a void as it were (at least with respect to *your* opinions), the onus for describing the differences falls to

you, the reader. Ergo, a prerequisite to criticism is to understand the opinions of the authors *as they were understood by them*.

Only now do we come to the third step, the third question, in careful reading.

(3) What do I think the results mean?

A complete answer to this question involves aspects of scope (restriction, expansion), credulity (do I believe the results? was there fraud?), rationality (were assumptions acceptable? is that argument flawed?), competence (were the measurements properly made? were statistical analyses correctly performed?), and alternative explanations (could something else explain these findings?). Notice that the answers to many of these parenthetical questions are not universal, but pragmatic. They depend on context, the problems to which the results were applied, to which we have turned to the authors for direction in step (2). Thus, for example, the validity of assumptions (for instance, about experimental methods, model equations, or statistical techniques) depends on the uses to which they were put--not necessarily their match to truth (although, of course, we aim in the end to learn the truth of the matter).

In conclusion, there are three rules for leading a discussion:

1. Start with a summary
2. Plan for the the life cycle of the discussion
3. Ensure that disagreement is well founded

In good discussion, like in genius, there is only the smallest fraction of inspiration. The rest is work!