

How to Give a Good Paper Discussion*

James J. Choi
Yale University and NBER

September 19, 2022

Abstract: How do you give a good paper discussion at a conference? I cover seven common pitfalls to avoid: (1) thinking your primary audience is the paper’s authors; (2) spending too much time summarizing the paper; (3) trying to squeeze too much into the discussion; (4) being too eager to present your own research; (5) not preparing enough; (6) being too negative in tone; (7) overestimating your command of the paper. I conclude with a non-exhaustive typology of useful comments in a discussion.

At some point in your academic career, you will probably be asked to give a public discussion of somebody else’s paper at a conference.

Giving a good discussion is an act of intellectual generosity, but it is also in your professional self-interest. Everybody is busy, which means that few people will read your papers. Even when they do read your papers, in this age where most papers are coauthored, the signals your papers contain about your research ability are commingled with signals about your coauthors’ ability. A discussion is a rare opportunity for you to grab the attention of a group of people who probably won’t read your papers and demonstrate your insightfulness and knowledge without the obscuring effect of coauthors.

The usual format of a discussion is to summarize the paper and then offer comments, criticisms, and suggestions. You will typically have 10-15 minutes to speak. Giving a good discussion is its own art—somewhat different than giving a good seminar talk. Unfortunately, we give almost no training to our Ph.D. students on how to discuss papers, leaving junior scholars to figure it out on their own (without much frank and honest feedback) once they start giving discussions. This document is my attempt to ameliorate some of this deficit by addressing the most common pitfalls in giving discussions.

* I thank, but do not implicate, Kelly Shue and Paul Tetlock for helpful feedback. Pedersen (2021) inspired portions of this document.

Pitfall #1: Thinking your primary audience is the paper's authors

Your primary audience is the conference attendees, not the paper's authors. The number of non-authors listening to your discussion always vastly exceeds the number of the paper's authors. It doesn't make sense to drag the entire room through points that aren't of interest to more than 95 percent of those present. Your discussion should help the non-authors in attendance understand the paper and its strengths and weaknesses better. An ideal scenario (which seldom occurs) is one where your discussion shapes the entire profession's collective assessment and interpretation of the paper.

Common manifestations of the mistaken belief that the authors are the discussion's primary audience include granular comments on the writing in certain passages or of the design of certain exhibits, pointing out typos in the paper, or making comments without providing enough context for non-authors in the audience to be able to understand them. As Pedersen (2021) writes, "A discussion is *not* a referee report spoken out loud."

This is not to say that the paper's authors are not *an* audience of your discussion. Hopefully, your discussion will help them improve their paper. Good discussions often live in the intersection of what interests the conference audience and the authors. Good discussions sometimes cover material that interests only the conference audience. Good discussions never spend significant time on content that interests only the authors.

Pitfall #2: Spending too much time summarizing the paper

Too often, I see discussants chew up over half of their allotted time summarizing the paper. It is rare that the paper summary is a significant source of value-add in a discussion. The authors already know what the paper is about. The audience has just heard the paper presented to them. Everybody wants to hear your original thoughts about the paper. They don't want to hear the paper presented a second time!

Almost always, the summary of the paper should be the "elevator speech" for the paper, the sound bite that helps consolidate it in the audience's memory. I try to keep my paper summary to less than 60 seconds.

Pitfall #3: Trying to squeeze too much into a discussion

It is common wisdom that seminar audiences will only end up remembering one thing about any given paper. That's for an hour and a half talk! In a 10-15 minute discussion, your audience's capacity to absorb multiple points is even more limited.

I try to present no more than three comments in a single discussion. If you have 10 minutes to discuss a paper, that's less than three minutes per point if you leave 30-60 seconds to summarize the paper and 30-60 seconds to recap your points at the end. If you have 15 minutes, that's less than five minutes per point. That's not a lot of time to make an argument about any substantial point. My philosophy is that less is more. If you try to squeeze in six points in a rushed manner, you will end up communicating none of them effectively. Better to present three points really well. Getting down to this length often requires deletion of intellectually significant content that you would definitely include in a referee report.

Closely related sins: Briefly flashing slides onto the screen that the audience doesn't have time to read (why make the slide at all if the audience is not meant to absorb its content?) and introducing a lot of mathematical notation (an audience has a very limited ability to remember what your γ and ξ are supposed to stand for in general, and in particular during a 10-minute presentation).

Pitfall #4: Being too eager to present your own research

Chances are good that you were asked to discuss this paper because you've done related research. Chances are also good that you will be too eager to present your own research during your discussion. We all are biased towards overestimating the quality and relevance of our own research, and it's also easier to present something you already know a lot about instead of mustering an original thought about somebody else's paper.

Sometimes, it is entirely appropriate to talk about your own research in a discussion. But apply a *much higher* hurdle than you otherwise would when making the decision about whether to do this. Your paper should be super-helpful for the audience's understanding of the contribution of the paper you're discussing or for the authors to improve their paper.

Pitfall #5: Not preparing enough

A discussion is like a 100-meter dash: unlike in a longer event, a stumble in the beginning can derail the entire presentation. This is why I script and prepare a discussion more meticulously than I would a full seminar presentation. Optimal preparation methods will vary by person, but my own practice is to type out every word I plan to say during the discussion.¹ I won't stick exactly to this plan when I give the actual discussion. But figuring that I speak at about 150 words per minute, the script gives me a very good sense of whether I will be able to fit my content into the allotted time.

Pitfall #6: Being too negative in tone

It's easy to get carried away when you're criticizing a paper. It is rare that being maximally blunt and harsh promotes the advancement of knowledge more than an approach that comes alongside the authors as a helper and a friend. Nor does being extremely negative make you look smarter to the audience.

A wise colleague of mine once said that nearly everybody in the audience will forget what the discussant said within a couple of weeks, but the authors will remember forever. This is not to say that you shouldn't criticize the paper's shortcomings. But try to criticize in the way that you would criticize a close friend's work.

Pitfall #7: Overestimating your command of the paper

The authors have been thinking about their paper for months if not years. You have been thinking about their paper for a couple of days. Who do you think is more likely to be mistaken in their thinking about the paper? If you spent a few hours trying to replicate the paper's results and failed, who do you think is more likely to have made a coding error? Be humble in your criticism.

¹ Typing out your entire discussion might feel like a daunting amount of effort. But if you speak at 150 words per minute, a 10-minute discussion is 1500 words long, which is a little less than three single-spaced pages. I find it takes me less time to type out my talk, revising as I go, than to run through it orally repeatedly while revising content in my head and checking whether I'm staying within the time limit.

A non-exhaustive typology of useful comments in a discussion

Less commonly, I see discussions where it feels like the discussant is filibustering to run out the allotted time without making any intellectually substantive remarks about the paper. If you're having trouble thinking of what to talk about in your discussion, here are examples of useful things to include:

- An independent attempt to arrive at the paper's result, so that we get a triangulation as to whether the results are believable. This could be through original data analysis, a synthesis of the prior literature's results, a back-of-the-envelope calculation, or a proof of the main theoretical result in a simplified setting.
- A simpler or alternative way to understand the intuition behind the paper's result.
- A discussion of the implications of the paper's results. If we believe the results, then how does that change the way we think about the world? (Is the implication not believable, for example because it requires accepting implausible effect sizes?) About prior research? Do the results suggest new model formulations? Are there certain settings in which we think the results are more likely to be true than others?
- Placing the paper's results in the context of prior literature. What makes the paper's contribution significant? Where does it accord or clash with prior work? Is there prior work whose overlap with the current paper diminishes its marginal contribution?
- An alternative interpretation of the paper's results. The authors say that the paper's results imply X, but you discuss how they could actually imply Y. The authors argue that a certain mechanism is responsible for their result, but you argue that another mechanism could be the primary one, or that the authors' mechanism is unlikely to be able to produce the result.
- A criticism of an identifying assumption in an empirical analysis, or a simplification in the setting needed to obtain the theoretical result.
- Pointing out a crucial potential confound in the analysis. This should not be a nitpicky request for another robustness check, but something that has a good chance of overturning the main result.
- Identification of an error or fragility in the methodology. Was an estimation procedure inappropriately applied? Were the standard errors incorrectly computed?

Was random assignment done improperly? Are the coefficients being driven by a few outliers? Is there a mistake in the proof? Does changing the theoretical setup slightly flip the results? Again, restrict yourself to discussing major issues that are likely to overturn the main result, not minor flaws.

- An explanation of the paper's methodology. Usually this type of discussion is not useful for the audience because it rehashes what was in the main presentation. But if the methodology is unfamiliar to most of the audience, broadly applicable, and you anticipate that the author's presentation will not have time to explain the methodology, this can be an appropriate type of comment. For example, I recently gave an introduction to neural networks and deep reinforcement learning in a discussion.

In general, discussions should stick to big-picture points rather than minor details that won't affect the major contribution of the paper. Also, remember that criticism that merely tears down a work is relatively easy. The ability to give constructive criticism that offers a concrete way to improve the paper separates the masters from the amateurs.

Reference

Pedersen, Lasse, 2021. "How to Succeed in Academia or Have Fun Trying."
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3972340