

The NSF reviewing process

The process

The typical NSF panel I've seen is composed of about 1/4-1/3 people who have worked in the area at one point in time, but may not be active/publishing, about 1/4-1/3 people who are in related fields, and the rest actively working in at least the broad area. Some program managers are better at recruiting people, and you may get lucky and have everyone be an active researcher. You probably have about a 30% chance of having someone who really knows your specific area read your proposal. You'll probably get one person who at least keeps up on the research in your area, even if they don't work in it themselves, and one person who knows very little about your area. On rare occasions the program manager will get an external reviewer (who isn't on the panel) to review it - this usually happens when your area is a bit further afield than normal.

Each proposal is (usually) read by 3 or 4 people. Sometimes it will be read by more for one or more reasons:

- Someone on the panel decided to read it, just because it looked interesting from the title/abstract, even though it wasn't assigned to them.
- Someone on the panel (on the day of the meeting) volunteers to read it. This can happen because of a conflict of opinion amongst the assigned reviewers, or because, while listening to the discussion, they decided they might have an opinion.
- If your proposal crosses disciplines, it will typically have an extra reviewer or two from another area.
-

Each reviewer is suppose to read the proposals before the meeting, and enter comments. A typical review load is 5-15 proposals, and, while some reviewers try to read them all the night before (!) usually they'll be read over two weeks to a month or so.

The reviewers don't see each other's comments (typically) until the day of the meeting.

The day of the meeting, each proposal is discussed in turn, and the proposals categorized as Highly competitive (Really Good, very likely to be funded), Competitive (Good -we like them although they have some flaw, but probably won't be funded unless money falls out of the sky), Low competitive (the idea's good - try again after fixing major flaws) and Not Recommended for funding (reject - don't even bother re-writing it). (The ranking scheme may vary a bit - but usually the proposals are grouped as they're discussed into roughly these conceptual categories.) If your proposal ended up in the Really Good category (and typically only about 1/10 do) then you will be likely be funded, barring any unusual circumstances. These proposals are usually ranked by the reviewers with Excellences and the occasional Very Good. Of the proposals in the Good category, the top one or two might be funded, if congress is generous this year.

The program chair listens in on the discussion, and provides guidance on the procedure, but basically doesn't contribute to content discussion. He/she has the final say on what gets funded. I think they generally follow the advice of the panel, but as I've never checked to see if the proposals we ranked highly were funded, I can't tell you if this is indeed the case. I have

heard of at least one occasion where the program chair killed an excellently ranked proposal because "it wasn't in an area I was interested in" (he was retiring that year...). But (hopefully!) this sort of thing is rare.

Note: Some small proposals (especially from under-represented schools/states) may get funded even though they weren't in the top 1/10 (but they better be ranked well). This is because there's other pots of money to draw from in this case (e.g. EPSCOR states). And (I think) a small proposal might just slip in to fill in the chinks in the budget.

After the first day everyone comes back and the proposal ranking is gone over again with a fine-toothed comb, with some proposals shuffling up and down. Then everyone sits down and writes panel summaries which try (with mixed success) to capture the flavor of the oral discussions.

How this affects your proposal review.

- The reviewer will remember, at best, your research area, a basic idea of your proposed research plan, and any unusual collaborative efforts. But that's only if you make a concerted effort to clearly lay out those points in the first two pages.
- They will definitely remember if:
 - The writing was bad/confusing/full of errors.
 - You didn't do what you said you were going to do in the first two pages.
 - You missed out someone else's research you should have been aware of.
- You should do your best to make sure they remember:
 - That you did an excellent job describing the current state of research in your area, and what the open problems are.
 - That you clearly laid out a list of problems and potential solutions.
 - Moreover, you included a summary of those problems in an easy-to-find place so the reviewer can find it again when needed.
 - You thought about the broader impact and education components.
 - (If applicable) That you have thought out how to evaluate your research, and/or apply it to real problems.

How the composition of the panel affects your ranking

If someone on the panel works in your area, this can be either a Good thing, or a Bad thing. But it will definitely affect your ranking, because the remaining panelists will defer to the person who "knows what they're talking about".

On the plus side: If you really do have a good idea, but maybe you haven't presented it as clearly as you might, or argued why this is an important research topic, a knowledgeable reviewer can help.

On the down side: If the reviewer doesn't think much of your idea, they can kill it. The reviewer's arguments will be one of the following: "I/someone tried this x years ago and it failed", or they may present technical reasons, or they may say you didn't cite some relevant work. To combat this, I suggest the following:

- If you've got demonstrable results, include them and, more importantly, make sure they're easy to find/stand out from the body of the text.
- Make sure your "why this research is important" and "why this approach will succeed" sections are clearly understandable to someone who is (presumably) reasonably intelligent, but not versed in your area.
- Include enough detail to convince someone in your area that you know this material cold, and that your solutions will work. The real reason for this section is so that the semi-knowledgeable reviewer (who likes your

proposal) can point to that section and say "this looks really convincing - explain to me why this isn't convincing you".

People outside of your area will primarily judge your proposal on the first two pages, the outreach section, and on how well the proposal is written (does it flow, are the non-technical arguments convincing). I have seen a proposal that was very well-written, but somewhat slim on technical merit, get funded because the three reviewers were not knowledgeable in the area. And, in general, a well-written proposal will get higher marks from less-knowledgeable people simply because that's the only thing they have to judge the proposal on.

Of course, this means that if your proposal is dense, full of jargon, and lacks compelling, easily understood arguments, it will get lower marks. If you're working in an area where the problems are not so obvious (e.g., if you work with robots, you know it is stunningly difficult to get a robot to roll down a corridor, but the average computer science researcher doesn't think this is a hard problem because they walk down corridors all the time, so how hard can it be...) then you're going to have to spend some time convincing the reviewer that this is actually a problem, and that the current solutions are not good enough.

Why proposals aren't funded

The following are the most common rejections I've seen. They come in pairs, more or less, with a good proposal balancing between the two.

- **It's already been done.**

- This usually happens with a proposal where the PI is new to the field, and just plain missed an entire body of research (possibly because it came under a different name in another field). Usually the review will be a kindly one, and simply say "go look at this work".
- It also happens (sadly enough) to new proposal writers because they didn't do a literature search. Please, do your homework.
- Sometimes this arises because of a mis-understanding of the proposed research, and a mis-remembering of what's been done. The best way to combat this is to make sure you've touched on all related and are explicit about what you're doing that they didn't.

- **Not enough detail/vagueness**

- This usually arises because the PI just plain hasn't figured out how they're going to proceed. If the reviewers liked the problem area and the approach, then this proposal is a very good candidate for thinking/working on some more, and re-submitting.
- You may argue that, if you knew how to do something, you wouldn't be asking for money to do it.. Don't confuse not knowing how to do something with not enough detail. For example, I don't know what roads to take to get from St. Louis to New Orleans, but I do know enough to pop-up Google Maps and get driving directions, and that there's probably a north-south freeway that parallels the Mississippi, and that it will take a day or two to get there.
- Be concrete about the problem statement and evaluation, even if you're not quite sure of all of the details on the approach. A poorly detailed problem statement is death to a proposal.

- **There's not enough research/nothing new**

- This is slightly different than it's already been done. I usually see this with research ideas that have reached the point where most of the things left to do are fairly straight-forward, or just require an application of existing, known techniques. Yes, engineering/follow through is important, but it's not what NSF funds.
- There's usually no fixing this sort of proposal. It may, however, make a great start/first third of your next proposal.

- **Too ambitious**

- I know of several cases where the research in a proposal that got rejected as "too ambitious" was mostly completed in the following year... I think the way to avoid this review is to include at least a few pages of fairly short-term goals that are clearly achievable. Then you can be pretty ambitious in the rest of the proposal.
- You may have oh-so-many ways to tackle the problem, and there's so many different ways you could use your work... that's GREAT, give some feel for the breadth of problems/approaches out there, then PARE the problems/approaches down to something that matches the amount of work you can do in 3-5 years.
- **Lack of evaluation/no application**
 - This happens if the reviewer thinks you're solving a non-existent problem, or if your proposed solution won't be any better than existing solutions.
 - If your work is just too far from being usable (yet) in a real-world application, at least come up with compelling test cases and evaluation criteria.
 - Think about how to evaluate your solutions/approaches.
- **Poorly written**
 - There are services you can hire that will help you clean up your writing. If you aren't a native speaker, or you don't write well, use them. Most universities have a writing center as well.
 - Write the proposal early, let it sit, and go back to it.
 - Don't cobble together bits and pieces from papers and expect it to flow.
 - If you're working with multiple PIs, designate someone as the coordinator and make sure everyone's reading everyone else's part.
 - Get other people to read it.
- **Technical flaws/panel doesn't think this will work**
 - It could be that you didn't explain your approach well enough, but more likely you're trying an approach to an existing hard problem that isn't likely to be solved with your (usually naive) approach. The only fix for this is preliminary data that shows your approach will work, or a re-working of the problem statement to bite off a smaller chunk that you're more likely to succeed in doing.
- **Just not compelling/exciting**
 - This is a hard one. A proposal can be well-written, have a solid technical approach, and still not excite anyone. The best fix for this is to find a compelling broader application (or a novel technical approach, but application is usually easier).
- **How to make a reviewer unhappy**
 - Don't spell check. Change formatting halfway through the paper. Use really small text. Don't check your figure numbers and captions. Use equations without defining the variables. Bounce around from subject to subject. Claim you're going to do x then actually describe y. List all of your achievements and describe all of your current research and spend one page describing new work. Define 20 equations without any accompanying text. Make it really difficult to dig out the broader impact, intellectual merit, and proposed work. Ask for tons of money for lots of students to do very little work. Assume your reviewer knows your field intimately, and jump straight to the details. Cut and paste three existing proposals from different people into one proposal, and add a summary page that "glues" the result together.

