
## THE PHD METAGAME

1. **Your Paper Is an Ad**
2. Don't Try to Reform Science
3. How to Pick Your PhD School
4. Don't Make Things Actually Work
5. How to Get Your Paper Accepted
6. The Cursed Word "Interesting"
7. Your Advisor Has Five Impossible Jobs
8. Try Even a Little at Conferences

Jan 10, 2025

## THE PHD METAGAME

# Your Paper Is an Ad

Insiders to our social practice of science understand that a research paper serves several purposes. In increasing importance, a paper is:

1. Currency
2. An advertisement
3. Brand marketing

These are a paper's lesser, primary, and greater purposes. To play the scientific metagame well, you must master all three.

This view is in contrast to what outsiders and newcomers to science believe, which is that the purpose of a paper is to share a novel discovery with the world in a detailed report.

## Your Paper is Currency

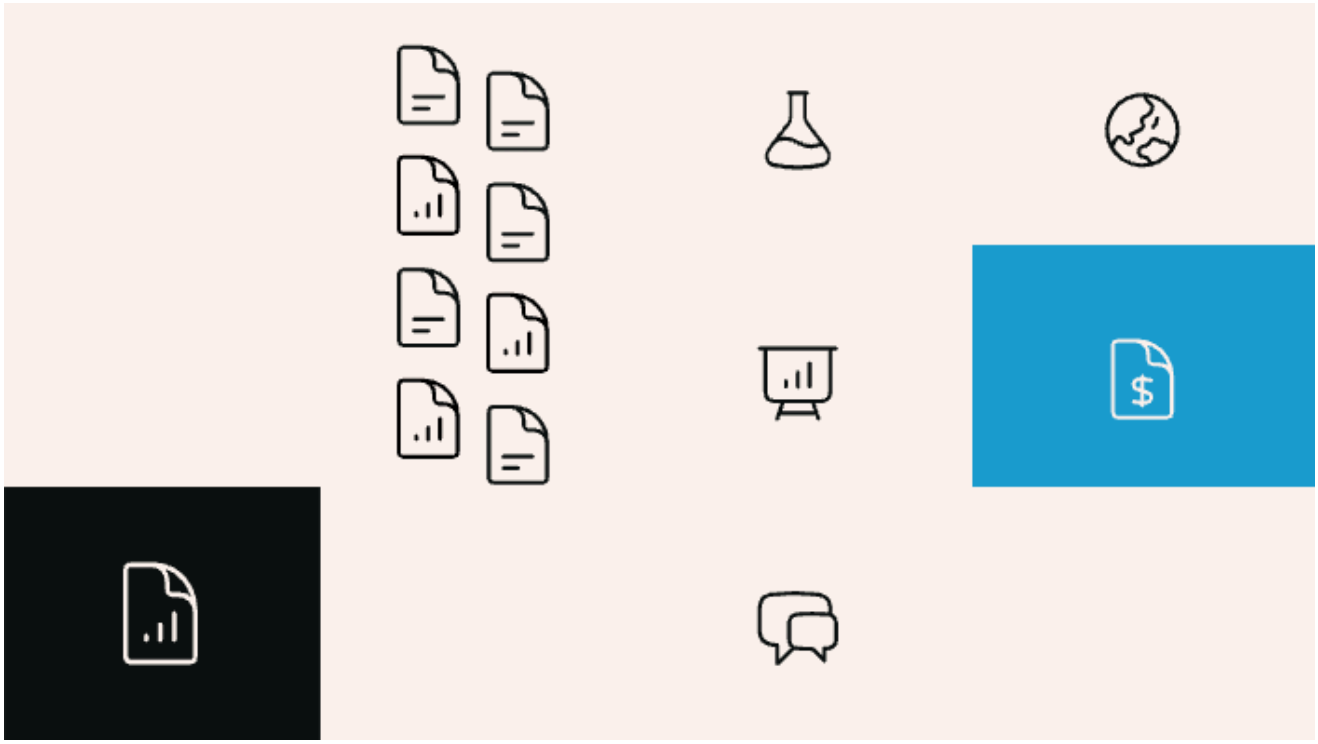

The lesser purpose of a research paper is as **currency**. Simply having "a paper" shows that you and your professor are active. It is a ticket for you

to go to a conference, which is important for socializing and building the human bonds that are the backbone of how we conduct science. And, as a bonus, you end up learning a lot just by doing it.

Because papers are raw research currency, you must build in yourself the machine of writing papers. This means finishing a paper's worth of work (the actual research), creating the artifact itself (text and figures), and getting it accepted. You must train such a capacity because you need to be able to produce many papers over the course of your career. It also builds confidence, a vital resource.

Furthermore, like shipping a product, finishing a paper teaches you an enormous amount about the process. For example, here are some things that you only observe in the back half of a project:

- Which statistics and narrative framing to apply
- What gets written about vs omitted in the paper
- How your advisor writes (you can see the shadow of readers here)
- All of reviews and rebuttals
- Conferences themselves
- Preparing and giving a talk[01]

If your paper gets ignored, it still works as currency.

At this point, we can already observe a particular failure mode among researchers: *only* thinking of papers as currency.[02] Such researchers get stuck in the local optimum of pumping out incremental improvements—quantity over quality. They never advance to targeting the deeper purposes of papers, such as creating intellectual advertisements.

# Your Paper Is an Ad

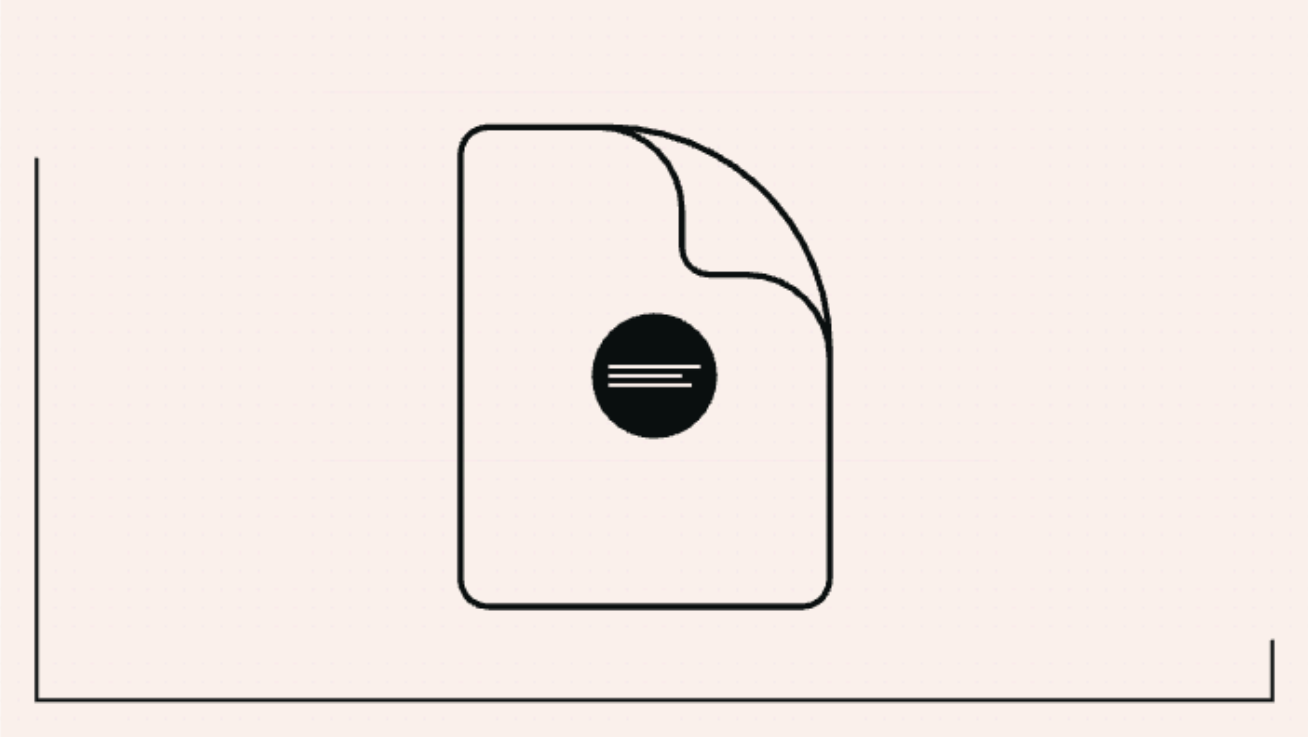

The primary purpose of a research paper is to advertise an idea.

It may seem derogatory to refer to research papers as advertisements. Not so! As we'll see below, a key reason why they must function as advertisements is that science is a collaborative process. The advertisement quite literally says, "hey, this idea is interesting, come work on it." If the ad worked, teams of researchers will verify, extend, break, or generalize the idea, yielding further scientific progress.

The requirements of an effective paper-as-advertisement are:

- One idea
- Pitched well
- Around a constructed central narrative
- For professors
- To achieve recognition

Knowing your paper is an ad is immensely clarifying. Most critically, it helps center you when you inevitably become sidetracked writing it. Because you are always zoomed in so close to your research, you can

produce countless pages of elaboration about your results. Understanding the requirements of a good ad helps you zoom back out.

## Put One Idea in Your Paper

People cannot retain more than the major idea per paper. This is why it is an ad for one idea. People are reading too many papers to understand an elaborate, nuanced, multi-idea novella. Plus, you only have eight pages.[03]

I must point out: do not take this "one idea" thing too literally. Your idea can have three parts, or whatever. Your advisor will know what to do here. But readers should feel the paper has a single direction. A guiding focus. A clear thrust. And if you have multiple significant ideas, just write multiple papers.

Think of it this way: a reader only has the cognitive capacity to understand a few points in a paper. It's like an insight budget. Your job is to make sure all of your insight budget is spent on points that make your paper as compelling as possible. Ideally, the points all revolve around your central idea.

## Pitch Your Idea Well

Your paper must be a good pitch for your idea because it cannot be the opposite: an comprehensive treatise.

Here is my claim: if an idea is sufficiently interesting, one eight-page paper cannot explore it exhaustively. This is why it is an *advertisement* for the idea, i.e., a good pitch.

Part of its ad nature is self-inflicted. To make a clear and compelling narrative in eight pages, you must omit myriad details.

But the other reason a paper cannot explore an idea exhaustively is the simple reality of scope. It takes multiple teams of people many tries to really understand anything. If one paper tried to comprehensively solve a research direction, it would either (a) never come out, or (b) be too large when it did come out. This reality is actually fine, because that is why we have science: many people can work on a topic over time.[04]

The best papers so compellingly advertise an idea that many people jump to work on the direction it proposes. Often, the original paper's conclusion turns out to be wrong, but the idea is good enough that it does not matter because it sparked a direction. This point is so important I am going to write it again. *Often, the original paper's conclusion turns out to be wrong, but the idea is good enough that it does not matter because it sparked a direction.*[05] See, for example, the flurry of papers around the The Lottery Ticket Hypothesis ↗.

In other words, a paper should be an ad that causes others to get excited about a research direction, and maybe even pursue it themselves.

## The Shape: A Constructed Central Narrative

Your task in writing a paper is to craft a compelling central narrative. The crux of the narrative is your idea. To support the narrative, you will provide necessary evidence from your experiments. Your target audience is both someone in your general field (e.g., NLP[06]), as well as someone working in your specific domain or task. (Your specific audience is professors. More on that next.) You must convince them both that your idea is interesting.

It's worth noting that this narrative isn't reflected in the paper's structure. Papers have a fixed structure, and deviations are usually unhelpful.[07] Here's the structure,[08] and where your narrative should show up:

1. **Title** ← Here
2. *Figure 1* ← Here, if possible
3. **Abstract** ← Here
4. **Intro ← Here, at length**
5. *Approach* ← Here, maybe / implicitly
6. Method
7. Model
8. Experiments
9. Results
10. *Discussion* ← Here, a bit
11. Related Work
12. Conclusion ← Interestingly, less-so here.[09]

This is why the introduction is the hardest to write, and is written last, usually by the advisor.

Constructing a compelling central narrative is challenging because at first because you did not realize it was happening. Before you write your first paper, you have probably read anywhere between ten and one hundred of them. Each had a central narrative that expressed an idea. So it seemed that a clear, idea-wrapping story falls naturally out of doing research. Not so.

What you actually end up with doing research is several inklings, a trail of mediocre work, and finally large a pile of experiments.[10] Usually, one or two experiments in the pile went well and triggered the writeup. The story that you tell is not the story of how you got there—i.e., the timeline of your research. You tell the idea's story.[11]

The idea's story often uses a shockingly small portion of your results. Instead, making a good idea narrative may spawn a whole new set of work required to flesh it out, such as:

- experiments with variations of your idea
- ablations
- skimming a dozen papers that are now relevant
- comparisons to past work
- human evaluations
- statistical significance
- the hard, slow, nebulous work of assembling the paper's narrative at all scales
- and, of course, all the normal paper-writing work.

All this can take surprisingly long. To make matters worse, when you get your key results, it feels like you are done. People will even say, "now we just write this up." Ha.

The other side of making a central narrative is you must purposely omit from your paper mountains of details that would obscure it: insights gleaned while playing with data, negative results, caveats, and weaknesses. You must cut ruthlessly to keep the paper clear and compelling. Emphasize cohesiveness in place of extensibility.

While writing your paper, such omissions will feel like a crime, especially because these topics are the most fruitful grounds for future projects. And, just as shadows are vital in understanding 3D form, limitations are vital in understanding the true shape of your work. But, you must remember that the goal of your paper is not to describe the true shape of your work. It is to create an effective ad.

## Write Mostly to Professors

A paper is a funny thing in that you are balancing two competing audiences:

1. Professors must find your idea interesting
2. Reviewers must accept your paper

Broadly, professors don't care about the details. Reviewers nitpick the details and ask for every metric under the sun.

This isn't a paper-writing guide, but I'll mention briefly here one strategy for appeasing reviewers: simply put in the paper everything they might ask for. This is a lot of work, but in my experience, it saves you time during rebuttals.

But sometimes the objectives are aligned. Reviewers who know what they're doing operate as follows: they make a gut decision about your paper, and then look for evidence to support their choice. This sounds crazy, but it's actually a good thing, because the question of "is your idea interesting?" is completely unquantifiable. So we use our guts.

It's obvious why you must appease reviewers: your paper must be accepted to a conference so it counts as currency. But why is entertaining professors the other objective?

Professors are the main actors of academia. They are the ones who organize conferences, can admit you to their university as faculty, serve on your committee, advocate for you, hire you as a postdoc, write you letters of recommendation, and spread word of your work.[12]

As such, (most) professors don't have time to read your equations, code, or appendices. They're operating in the space of ideas.

Don't worry about other grad students. If they need to know more gory details, they'll email you.[13] Ideally, release your code and data to help

them out.

## Your Paper Wants Recognition

The goal of your ad is to be recognized for your idea. Recognition is the fundamental currency in science. Or, put more bluntly, fame. You must be known for what you do. This is not as dystopian as it seems, and it is—once again—because we practice science collectively as a large group of humans, rather than conducting abstract "pure" science in isolation.

So, the goal of your ad is the goal of all ads. Occupy headspace. Your idea must stick in people's brains, but also *you* must stick in people's brains.

When you are starting out, it is mostly your idea that people remember with a good ad. But for seasoned advisors, it is the other way around: "Oh, X's lab did another paper on topic Y."[14]

Recognition is why long author lists are a wash for middle authors. Social credit gets attributed to first author and last author. If there are three authors, and you are the middle one, I think you can kind of remind people all the time that you were a close contributor to the paper. But beyond three, every author in position *2* through *n-1* is just getting their CV filled out without any direct ad benefit.

So, along with granting research currency, we've covered papers as ads: one idea, pitched well, with a constructed central narrative, primarily for professors, to achieve recognition.

But while a paper is an ad for an idea on the surface, its true purpose lies in a larger brand marketing strategy.

## Your Paper is Brand Marketing

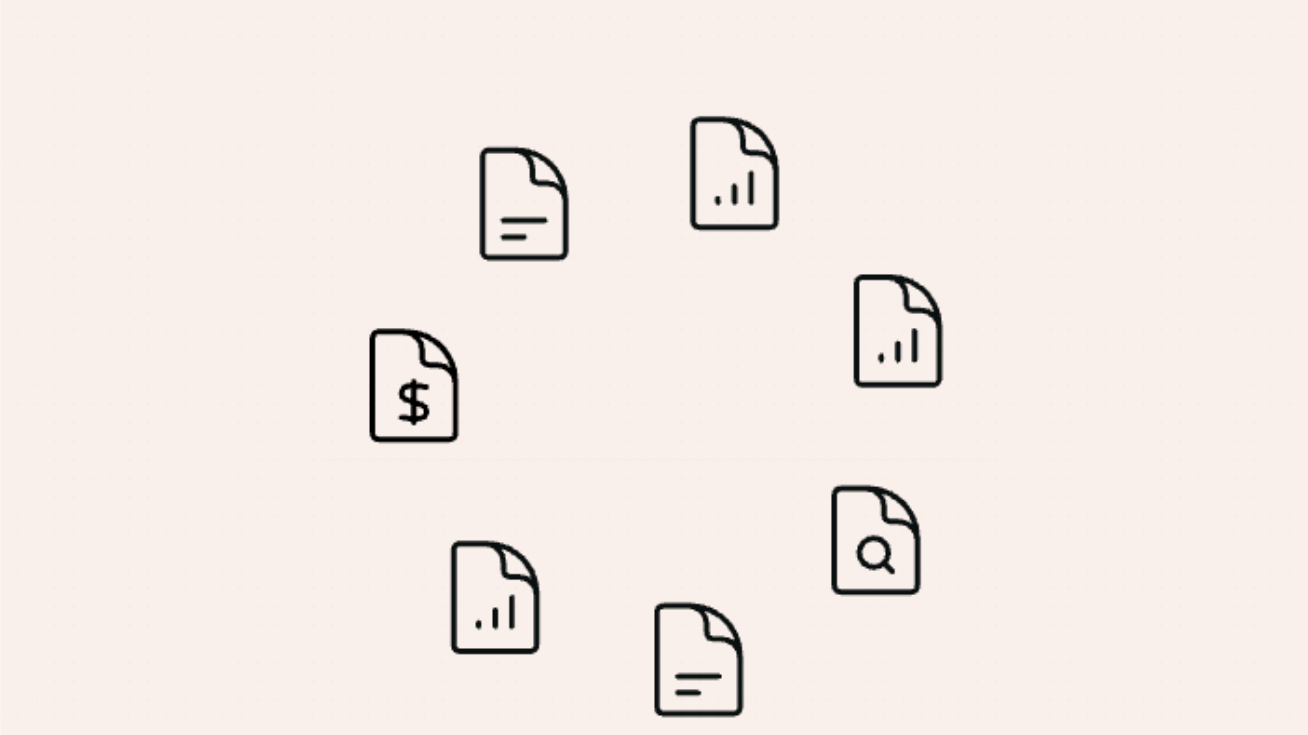

The purpose of a research paper is not to advertise an idea to your scientific community. That is a *requirement* of a paper. If a paper does not clearly and convincingly do that, it will be rejected or ignored, which are effectively the same thing.

If a paper is actually noticed, then it becomes one component of a brand marketing strategy. This is both your advisor's brand and—eventually—your own brand.

Brands are important because people need something easy to remember about you. It is not reductive. It is because you are in a field with (tens of) thousands of people, and not everyone can remember that you love dachshunds or your favorite flower is a lily. Researchers, labs, and institutions want to grow their reputation and establish name recognition for certain scientific ideas and schools of thought. Some examples: Bayesian enthusiasts, social activists, people-who-work-on-specific-task-X,[15] task-makers, or my own former lab's focus on common-sense reasoning.

A way to see why brands are important is to see what happens when you do not have one. Someone asks you the simple question, "what do you work on?" Then, you fumble and say, "oh, well I did a bit of X, then I worked on Y, and I also am interested in Z right now." Each of those were probably minor contributions because it takes time to get into an area. This person leaves the conversation not remembering X / Y / Z, but with an impression that you haven't quite found your niche yet.

The fact that your paper markets a brand causes otherwise inexplicably weird things to appear in it (your paper). For example, if your advisor likes Bayesian stuff but you don't really care either way, you will notice that the paper still talks about Bayesian stuff a lot. In fact it kind of seems to be the whole central narrative, doesn't it? You didn't think it mattered that much, and in fact the other methods were a lot easier to understand and sometimes seemed to work better. But if your advisor has a Bayesian brand, then you probably do not want to fight it. They have spent a lot of time and effort establishing their brand, and it is important to them that their brand continues.[16]

Brands take a while to establish. They do come with benefits, like a buffet of genuine insights and battle scars from publishing several papers in one niche. This is why professors resist field-swinging advances for so long: they agonize—perhaps subconsciously—at the thought of pivoting their

hard-earned brand, specialist knowledge, and membership in a sub-community.

When you start out, it's too much to worry about brands. But knowing about them will help you understand where your paper is headed. And if you're choosing which professor to work with, knowing their brand in advance will prepare you for what kinds of research you're likely to pursue.[17]

# Conclusion

This post describes a very specific lens through which you can view academic papers: as advertisements. These ads seek recognition, the fundamental currency of our social practice of science. The recognition is for both for an idea (surface) and for you (deeper). Any published paper has the base value of a CV line item and conference ticket. But by producing effective ads—by crafting a compelling narrative and omitting peripheral details—you unlock their true power, which is marketing an academic brand.

The game of science remains unchanged: seek truth via the scientific method and share it with others. As workers assembling the line-item facts that will be entered into the annals of history, it is vital to stay faithful to this ideal, even if it becomes obscured by thick layers of agenda.

But as a PhD student, you are also playing the PhD game. And the best strategy of the PhD metagame is to write several well-received papers. It is important to be aware of this angle of how papers function. Even if you only come away from this post with a vague awareness of these hidden objectives, I think you'll be better off.

## FOOTNOTES

01 Hopefully. Might be a a poster though. ↩

02 I've seen firsthand the failure mode of producing a stream of pointless currency-only papers happen to undergrads, PhD students, and postdocs. These are papers with an idea is so small that it could be a single sentence: *"We applied existing technique X to task Y, and got Z% benchmark improvement."* Unfortunately, there is no shortage of such papers. They disguise their true nature by filling eight pages and masquerading along with work of vastly greater novelty and interest. I do think it is probably OK to write a couple of these to collect currency and momentum if you need to. ↩

03 Please mentally substitute in your own particular constraints throughout (e.g., six pages). ↩

04 Having multiple independent teams tests an idea's robustness, too. That we can do this is a crowning glory of the scientific process. ↩

05 Depending on where you are in your PhD, this fact might blow your mind. (It kind of blew my mind even just writing it down.) Because for the first time, you will separate the notion of "conclusion" and "idea." The idea of your paper is broader and more nebulous than the conclusion you got from your experiments. The idea has the flavor of, *"what if we looked at this kind of thing in this way?"* ↩

06 Natural Language Processing, my home field in computer science. ↩

07 Structural novelty is usually unhelpful because you're burning a bunch of the reader's cognitive budget on your paper's roadmap. ↩

08 This is exhaustive. Sections are often merged. ↩

09 One funny trick: Conclusions are often the literal description of what the paper actually did. If you want to know the meat'n'bones of a paper's actual technical achievement, a great hack is to read the conclusion first rather than the intro. ↩

10 This characterization of the research process is a complete lie. In the majority of research I've seen (perhaps 90%+), some conclusion is already being aimed for at the start of the project, and the work is getting the details to align. Furthermore, depending on the nature of your project, the gruntwork may not be experiments, but proofs, human studies, or data collection. There's also all the nebulous work of figuring out what the project is in the first place. ↩

11 The difference between your research process and the idea's story is like the difference between famous scientists' notebooks and textbooks. One records the organic process of discovery, whereas the other presents the idea in idealized logical scaffolding. ↩

12 Plus I'm sure a lot more stuff. They seem awfully busy. ↩

13 Aside tip: email other grad students. They are also toiling for years for slivers of recognition. And they don't get much email. Saying hi and sharing a kind word about someone's work goes a long way. ↩

14 If the whole situation is starting to sound weird, it's because it is. When I graduated from my PhD, the sense I got from talking to friends and family was they believed I was now an expert at something. What I couldn't articulate was, "Actually, I just got OK at making ads for academics." As usual, a surprising amount of time was spent not realizing that's what I was doing. ↩

15 Back in my day, we had different camps of people who each specialized in a single NLP task: syntactic parsing, semantic parsing, machine translation, named entity recognition, language modeling, etc. ↩

16 The paper ↔ brand connection also causes events earlier than writing your paper, such as the research topics you work on in the first place. ↩

17 Since this is the deepest layer we're covering, you may think: are there deeper paper objectives than academic brands? I think the answer is yes, but they're more akin to culture in that they're not crafted explicitly. Funding agencies may put forward research themes or challenges that many groups work on, but I think these usually don't affect the course of science much because (a) professors still primarily value their brands, (b) the government is never uniquely clued in to the future of science in a way researchers aren't. The bigger trends come from what works. In NLP, these were the eras of rule-based, statistical, deep-learning, and now foundation models. ↩

## THE PHD METAGAME SERIES

NEXT →
2. Don't Try to Reform Science

## POST INFO

**THANKS** to Max Cerami for reading drafts of this.

**PUBLISHED** Jan 10, 2025

**TAGS** research

**INBOUND** ⇠ Your Advisor Has Five Impossible Jobs