

The Death of Peer Review

How the internet rebuilt theoretical science

Brian Wecht

The study of particle physics is nothing less than an attempt to understand nature at its smallest scales by describing the constituents of all matter. To be a particle physicist is to be a reductionist. Current experimental evidence indicates that our universe is made up of many fundamental particles, such as quarks and electrons, but it is currently an open question whether or not these can be themselves understood as avatars of a yet-to-be understood physical theory. Some people believe strongly in the possibility that what we see as particles may in fact be excitations of very small strings. These people are called string theorists, and this is their story. Sort of.

I am a postdoctoral researcher in particle physics and string theory at the Institute for Advanced Study in Princeton, New Jersey. At the current stage of the discipline, modern theoretical particle physics is very similar to abstract mathematics, in that good work must satisfy two basic criteria: It is new, and it is correct. *New* here means what you think it does, that nobody has done it before. *Correct* means that the work is logically consistent, and does not conflict with known experimental data. This experimental requirement is a bit fluid at best, since some of the work currently being done has no obvious consequences for experiments any time soon. Lacking this, the goal is to be logically consistent, both with oneself and with the rest of the literature—at least until the big experiment that everyone is looking toward: the Large Hadron Collider (LHC) in Geneva.

If this sounds pointless to you, you're not alone. There are many people out there for whom string theory is emblematic of all that is wrong with modern theoretical science, as its results may not be currently verifiable.¹ Lack of verifiability doesn't mean lack of rigor, however: The life of a theorist consists of coming up with ideas, and then meticulously checking them through sometimes

¹ Or ever, but let's move on from that, since it is an entirely different, and difficult, discussion. Suffice it to say that this issue has been seriously

arduous calculations. Although there are certainly some flamboyant types in our midst, by and large, we are careful to say only what we can back up with mathematics. But as the research we do in contemporary theoretical science has become more competitive and more rapidly paced, it has had a dramatic effect on the scientific method—particularly the peer review process.²

The scientific method is, as any sensible person will tell you, the best process our society has found for illuminating the fundamental properties of nature. The core ideas most of us are taught in elementary school—hypothesize, experiment, interpret, reformulate—are still there even at the level of professional scientific practice. The process must be modified slightly to accommodate theoretical sciences, such as abstract mathematics or some branches of theoretical physics, but all modern active academic scientists are doing Science, and Science involves an equally important step that we are often not taught as children: publishing.

The traditional process of scientific publishing—submission, criticism, interpretation, revision—is a microcosm of the scientific method. Typically, this process can take many months, if not years,

pissing people off for a while now, and I don't really have anything to add to it here except to say that theoretical science is awesome and haters can just shut up. Moving on.

² A quick digression on terminology: I interchangeably use the terms *theoretical high-energy physics*, *theoretical particle physics*, and *particle theory*. The *high-energy* part of these comes from the fact that the structures such physicists study become relevant only when things (like protons, but really whatever) are smashed together at very high energies. Technically speaking, this could include strings, which are not particles (strings are one-dimensional, and particles are zero-dimensional). Thus, many people in my field who work on strings will describe themselves as high-energy theorists. Still, out of a sense of historical obligation, many such people would say that they're theoretical particle physicists. The point is, for my purposes here, they're all the same.

and can be as much a matter of finding an appropriate journal for your results as it is checking that the results are verifiable and correct. At the end of the day, though, the goal is pretty clear: find something new that helps us understand some aspect of our amazing universe, and get it out there for other scientists to see.

What you may not know is that, although this process of continual submission and reevaluation is ongoing in every branch of science, there are some, such as mine, in which it is completely and utterly beside the point. In these fields, nobody reads journals any more, nobody wants to wait months for results, and peer review has become purely a matter of consensus rather than the job of a handful of referees. In short, the traditional peer review process that you know is dead, and theoretical science could not be better off for it.

A little history first. In the olden days (read: before roughly 1995), pretty much everyone read about the latest results in journals. But the review process of a decent journal was typically a long one, extending months, if not years. For scientists on the cutting edge of research, this was simply too long to wait to see what their colleagues were doing. To avoid this delay, they typically sent out preprints of their papers to other scientists while the manuscript was worming its way through peer review and the journal's editorial process. In those days, if you were a scientist, it helped tremendously to be at a prestigious university, since you would then have access to the most recent preprints, if not seminars describing the most cutting-edge research.

The internet provided a natural simplification of this process. Why send a preprint by snail mail when you could instead send one electronically? In 1991, Paul Ginsparg—then at Los Alamos National Laboratory, currently a professor of physics and computing and information science at Cornell University—started up an e-mail distribution list of electronic preprints, called e-prints, for theoretical high-energy physics research. Over time, the list grew to include more areas of physics, and eventually became the arXiv (pronounced

archive; the X is shorthand for the Greek letter χ), a website today used by physicists, mathematicians, biologists, and other scientists to post their latest research.³

I think it's fair to say that when the arXiv first started (long before I was doing physics), e-prints were viewed pretty much on the same footing as preprints. They were a way to get your paper out to your colleagues quickly while you waited for the print journals to publish. But of course, once papers became available online, the incentive to read print journals became almost nonexistent: At the time I'm writing this (July 2009), arxiv.org has over half a million e-prints under fourteen broad subject areas, many of which have numerous sub-subjects. The math section alone has thirty-two different subcategories, which is kind of an anomalous example, since not every subject has subdivisions, but hopefully it gives you the idea that posting to the arXiv is not a fringe activity. For many areas of science, particularly theoretical particle physics, this website is the main vector for spreading new research.

Before we get into the effect that the arXiv has had on scientific practice, I should note that I'm far from the first person to point it out.⁴ Suffice it to say that a typical day in the life of a modern theoretical high-energy physicist features one unalterable event: checking

³ It's probably not worth it to delve into a detailed history of the ArXiv here, but if you're interested in checking out a truly excellent primary source, I highly recommend arxiv.org/new/91-94.html, which collects many of the e-mail updates from the ArXiv in its earliest incarnation. Ginsparg has a notoriously sharp sense of humor, which is more than evident (I'm being diplomatic here) in these updates. To quote one from January 1994: "N.B. these are 'archives,' not 'bulletin boards' (persons using the latter terminology will be excluded further access)." Here's another, from May 1992: "please do not send comments or complaints regarding this message, regardless of how manifestly obvious or thoroughly objectionable you find it." I could read these updates for hours and not get bored. N.B.: This is not what a theoretical physicist does all day.

the arXiv for the latest research papers. These papers appear every weekday at 12 A.M. Greenwich Mean Time (8 P.M. EST), and like many theoretical physicists in the United States, I check the arXiv every evening without fail. As I said, the arXiv encompasses many different subject areas, but in theoretical high-energy physics, there are three main contenders: hep-th (high-energy physics—theory), hep-ph (high-energy physics—phenomenology), and hep-lat (high-energy physics—lattice).⁵ Each of these subarchives has a unique character and focuses on a particular subfield of high-energy physics.

As particle physicists are usually a pretty intense bunch of folks and also extremely hard-working, near-instant access to the newest papers has changed the pace of research in high-energy theory dramatically. The introduction of the arXiv meant that the only thing preventing your idea from getting out into the world and wandering around was that it wasn't written up yet. It is now not unheard of for a paper to be posted, only to have response papers posted the next day.⁶

You might think that the high volume of papers—in hep-th alone, an average of six or so daily, up to more than a dozen on a heavy day—would mean that the journals would still fulfill an important

⁴ Many scientists have discussed the effect of the ArXiv on the peer review process, most notably Ginsparg himself. See Paul Ginsparg, "Can Peer Review Be Better Focused?" Available at people.ccmr.cornell.edu/~ginsparg/blurb/pg02pr.html.

⁵ Since the problems with high-energy physics happen at small-length scales, one way to make these problems manifest is to pretend the theory lives on a discrete grid, or lattice, of spacetime points some set distance apart from one another. As you take the lattice spacing smaller and smaller, you can then regulate or understand the resulting physics that emerges in the real, continuous spacetime. People who study this are called lattice gauge theorists, and hep-lat is their domain.

⁶ Sometimes these have great titles that take potshots at the original authors, but usually they're called something like "Comments on [other

function, by virtue of the filter of peer review. But since pretty much all theoretical physicists are skeptics and insist on verifying everything themselves anyway, the added value of the peer review process is essentially zero. And since you can instantly communicate with your friends and colleagues to see what they think of the work, it is possible to figure out what the consensus opinion is. So, gradually, the journals have become essentially irrelevant to the lives of theoretical physicists, and the only thing that matters is the arXiv.

Thus the peer review process as we knew it when journals roamed the Earth is dead. But in the consensus opinions that form (typically over a few days or even weeks) around the papers on the arXiv, peer review is, in another sense, very much alive. And since the pace of research has increased dramatically, the scientific community can get places far faster than it ever could have before. Sure, a lot of sub-par research gets posted to the arXiv, and much of it would not have made it through the review process. One could even argue that *most* of the papers on the arXiv are like this, although I am personally neither that pessimistic nor that willing to piss off my colleagues. But even so, the increased pace of quality research is far and away a net positive for the field, and for scientific progress in general.

For a layperson visiting the arXiv, the field might appear to be a hopeless mess. The arXiv contains papers that contradict other papers, or, in more extreme cases, themselves. On many topics, there's a lot of confusion about what's right and what's wrong. But this is always the case in science, and I don't think there's anything particularly special about this aspect of it because of the arXiv, except that the confusion can happen more quickly. Besides, generally speaking, there is a consensus on most matters; from talking to people, reading

paper].” I'd like to see a list of the snarkiest paper titles, but as far as I know nobody has made one yet. If you want to see some real vitriol, the abstracts are the place to look, though in fairness to my colleagues, I should say that this practice is quite rare (but still hilarious).

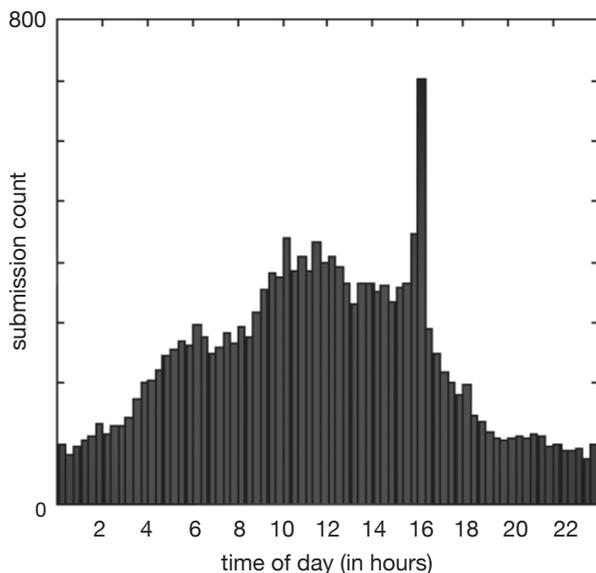
the papers, and checking the results, you get a feel for what is good work and what is not—and the fact that articles have not been officially peer reviewed matters not at all.

Of course, most theoretical physicists continue to submit to journals: Many official sources of funding ask to see peer reviewed articles, and young physicists who don't submit their articles to journals are effectively denying themselves tenure. This is in part because physicists in other fields do not operate as theoretical physicists do, and publishing in journals is thus an important part of their lives. As in other areas of academia, there are good journals and less good journals, and having a paper in a good journal can be a real feather in your cap when it comes time for promotion, grant approval, or whatever. My point is that, within the high-energy theory community, people publish in journals for reasons that have nothing to do with the actual research. The process of peer review happens organically through seminars and conversations, and the once-necessary formal structure has become irrelevant. And this is awesome.

The arXiv's effect on the scientific process is mirrored in a sociological effect on the scientists themselves. On a typical day, as I said above, you'll find six or seven papers on the hep-th section of the arXiv. These papers are displayed on the page sequentially so that you can see titles, authors, and abstracts. An interesting fact of life on the arXiv is that the ordering of papers is determined by the order of submission: Papers submitted earlier appear toward the top of the list, and to appear the next day, you must submit your paper by 4:00 P.M. Eastern Standard Time. If you submit your paper at 3:59 P.M. on Wednesday, it will appear with Thursday's releases. Submit at 4:01 P.M. on Wednesday, and you're up on Friday.

Where there's a sequence, there's jockeying for position. Add on a group of hypercompetitive workaholics, and you can bet there's going to be the occasional race for first. It is not at all uncommon for people to wait to submit their paper until immediately after 4:00

p.m., hoping that they'll appear first the following day. Here's a graph for astro-ph (astronomy and related stuff) of the average number of papers submitted during each hour of the day.⁷



Notice the insane spike at 4:00 P.M.

This kind of behavior is not at all uncommon, but it's pretty harmless. Or so I thought, until I saw that the same paper that released the above graph also demonstrated a correlation between

⁷ From Asif-ul Haque and Paul Ginsparg, "Positional Effects on Citation and Readership in arXiv," available at [arXiv.org/abs/0907.4740](https://arxiv.org/abs/0907.4740), to appear in the *Journal of the American Society for Information Science and Technology*.

the position of the paper and the number of citations: Papers higher on the list get cited more. Frankly, this should never have been made public, since it turns the submission race from a frivolity to something that may have consequences. Of course, at some level it's totally stupid: The thing that determines the number of citations is clearly the quality of the result (or famousness of the authors, if you're feeling cynical). Nonetheless, there's a measurable effect in being among the first to post—and to answer the question before it's asked, yes, I've done this. There is a small feeling of pride you get by clicking on your subject area and seeing your paper first. It's as if the arXiv looked through all the papers that day and determined that yours was the best.

A related issue is who gets to say that they discovered something first. Depending on whom you talk to, many high-energy theorists will say that results posted within one or two days of each other should be considered simultaneous (and are). But it's also fairly common that competing groups become aware of each other's efforts, and then arrange to have the papers appear on the same day. You might not think that there's anything like an emergency in theoretical physics, but you're wrong, and the arXiv is why.

Let's say that you find out that group X—whom you know/like/hate/admire—is working on a particular topic. In discussing physics with them, it comes out that your group's work is closely related to theirs. They tell you that their work will be web-ready in a week, although they're not willing to tell you their exact results. You arrange to post at the same time, but you're also playing it close to your chest, and won't divulge your own results. Now it is *on*. Your collaboration races into overdrive. Collections of notes, equations, and figures get thrown together in a draft. You stay up all night trying to derive something, only to find an error the next morning in the second line. After several frantic days, you post. Your career is saved.

This has happened to me twice, and although in a sense it's fun, in another, more accurate sense, it sucks. In one instance, my adviser and I were racing to have a draft posted before the 4:00 P.M.

deadline. He had the latest copy, and was in charge of posting it. This is when you hope to God that the page that tells you, literally, to take a deep breath and count to ten before posting won't find any errors with your submission (almost all submissions are done in TeX, a program that allows us to type up papers with complex equations in them. TeX needs to get compiled, and sometimes things can go horribly wrong). I was frantically trying to get my adviser on the phone ten minutes before the deadline, but with no luck. With two minutes to go, he finally picked up his phone, blurted "*I am literally freaking out*" and hung up. It was awesome. We got it in. Our results turned out to have almost no overlap with the other group's.

The other instance was not a planned simultaneous submission. Instead, a competing group (which happened to include a guy down the hall from me; don't ask) let it slip one morning that we were working on the same thing. I rallied my collaborators, and we agreed to contact him to see if he'd be willing to post at the same time we did. I went into his office to ask him this, and he told me he'd just posted. I am still mildly peeved about this.

A few paragraphs ago I mentioned the number of citations a paper gets; as if my field didn't need more ways of making people feel inadequate or underappreciated, this number can easily be found online. The website SPIRES⁸ is an invaluable tool for the workaday particle physicist, since it allows you to search the literature by author, date, journal, university, and other criteria. But it also keeps an index of how many citations a paper gets, and this number is updated every time there's a batch of new papers—that is, every weekday. Moreover, if you search your name and use the helpful citation summary sort, it will tell you how many of your papers are very well known (250–499 cites), known (10–49 cites), or, worst of all, unknown (0 cites). There's nothing more heartening than to see that your oeuvre consists entirely of unknown papers. And since

⁸ slac.stanford.edu/spires.

newer papers generally have fewer citations than older ones, there's always an effect where you're convinced that your recent work is substantially worse than your previous efforts. You can argue that people really shouldn't take these numbers too seriously, but this argument will fall on the deaf ears of a theoretical particle physicist frantically refreshing his citation-count webpage to see if anything has changed.

Another aspect of the internet age is the omnipresence of fad research. These waves happen once or twice a year, and go something like this. A paper appears and proposes a new idea, which opens the door to many generalizations or possible calculations. Particularly if the paper is by a famous person, people take notice, and get to work calculating. Within a week, the first follow-up papers appear. Note that I'm not implying that these are bad papers, just that they appear quickly. More and more people jump on the bandwagon, producing papers of various length, quality, and importance. After a while, the fad dies down, and the number of papers on the topic dwindles to a trickle.

Often the fads never really go away entirely, but stop becoming a main focus of the community. Occasionally, what begins as a fad becomes a crucial aspect of our understanding of high-energy physics. The fads of recent years include pp-waves, Horava-Lifschitz gravity, unparticles, pentaquarks, and M2-branes. I say these names mostly to amuse you. People who hate string theory—and there are a lot of them—view this faddishness as a negative thing for the field. They infer from it that string theorists have low standards, nothing better to do, and an inability to think for themselves. This attitude does the whole field a disservice. Sure, lots of crappy papers get written during these fads, but a lot of great papers get written, too. Claiming that string theorists are more prone to groupthink than any other collection of scientists is patently absurd; modern academia is full of such research fads, and although I don't have any numbers, I seriously doubt that it's any worse in string theory than

in any other discipline. Yes, there's certainly an element of time-wasting to it, but that's what modern particle physics is all about, at least until the LHC starts up.⁹ Once there's real data to work with, people will focus on that, and I suspect many of the fads will go away, even if careers have been made on them.

The internet has changed the lives of modern theoretical physicists in several other ways as well. Science blogs have achieved a prominent place in the community, most notably during the so-called string wars, in which some prominent bloggers took it upon themselves to criticize or laud string theory. These blogs had a marked effect, to the extent that public opinion of string theory (and by extension, string theorists) took a nosedive a few years ago. String

⁹ The LHC has been a bit stubborn, and had to be shut down almost immediately upon starting up because of problems with the magnets that direct the beam. This got a lot of media attention, but is really nothing more than the growing pains of a hugely complicated experiment. Further, there was some attention to the remote possibility that the LHC could create a small black hole. The probability that this will happen is exceedingly small, and even if it did, the black hole would probably Hawking-evaporate immediately, so there is categorically no danger of anything bad happening. Nevertheless, it is very difficult to convey the difference between 0.00000000000000000001 and 0 to the public, because the only thing people hear is the word nonzero. Let's establish this once and for all: Although the probability that the LHC will create a dangerous black hole is technically nonzero, it is so vanishingly small that you would have to collide things together for many, many lifetimes of the universe—14 billion years—to ever see it happen once. Nothing bad will happen, there is no Conspiracy of Mad Scientists, and the Earth remains safe from manmade black holes. There is a bigger chance that the Earth will be destroyed by Bigfoot, the Loch Ness Monster, the Illuminati, and the Greys coming together to form a Voltron-like super monster and eating the world's supply of magnesium than there is that the LHC will create a destructive black hole.

theorists didn't do themselves any favors, as they either remained silent, did not deign to comment, or said crazy things that made them look like jerks. Since then, many well-known scientists have started blogging, and it is not at all uncommon for results to be dissected by bloggers.

Recently there have also been some fun instances of data leaks—albeit not serious ones—propagated across the internet. One of the things that caused the recent spike of interest in dark matter¹⁰ among particle physicists was a slide of a preliminary photograph of some data shown at a conference presentation. A number of attendees at the conference took pictures of the slide and posted it to blogs, or circulated it among their friends, before the paper accompanying the presentation was posted. Many papers were written trying to explain an effect that existed in a preliminary photograph.

The fads and squabbles, however, are really just part of a much larger and extremely exciting dynamic in the way theoretical science is conducted. Though the conventional way of reaching scientific consensus has by now effectively disappeared in theoretical particle physics, the study itself is more robust than ever. Scientific research happens at a faster pace than ever before in the history of humankind, simply because having the internet at our disposal has so rapidly increased the pace at which ideas can be communicated over great distances. The resulting spirit of competition, although annoying at times, has resulted in a deluge of interesting new ideas, and in a generation of physicists working harder than ever at understanding our universe. With results from the LHC just around the corner,

¹⁰ We know via cosmological data that we see less matter than there should be in the universe. The difference between the matter we see and the matter we know should be there is called dark matter, which is just a catch-all term for "matter we can't see." The source of this matter remains completely mysterious, and is a major topic of work among physicists nowadays.

we can hope that we are on the verge of a sea change in the way we picture the physical world. Within five years, we may have either completely confirmed our present picture of the standard model of particle physics, or we may need to discard it entirely and think of something better. Either way, the discussion will progress apace via electronic preprints, as the community of physicists reaches an ever-closer approximation of the truth.