Many years ago, I was writing a project proposal – my first one that would go on to win funding. I had written one proposal before that and failed. This was a much smaller proposal, with a different funding agency, with a tradition of supporting first time applicants, but there was reason to worry. I was writing a proposal in an inter-disciplinary area that was new to me, with only a handful of people working in it worldwide, and I did not yet have a paper in that area, although one was in preparation. I did not know where the field would be in three years and worried that the problems I was proposing to investigate might not have solutions. So I went to a senior colleague with an illustrious record of successful grant applications and asked him how I could write a proposal when I was not sure if the project could be completed. He said, "Do not write about what you may or may not be able to do. Write about what you have done or almost done but have not yet published – write that as if you are proposing to do it. Once you win this project, write up your results and use the money for the next problem, which you will write about in the next project." I was shocked – surely it was not right to ask for money to do what I had already done! But I could not deny that the advice was practical and valuable. A collaborator who was from another country advised, "Do not worry about solving all the problems. Say what you wish to do – if in the end you do not succeed, so be it." I followed both suggestions to some extent – my proposal included the unpublished but nearly completed work as a goal and also some of the ideas that might not work out.

The proposal was funded, but it started me thinking about the funding process. Why are research grants needed, who needs grants at what level, how are grants decided, and most importantly, do research grants always lead to good research? Since then, I have had many discussions with my colleagues – mostly physicists – and have also taken part in the peer review process of funding agencies as well as the review process after funding. There seems to be no doubt among scientists that the processes associated with research grants – proposal, review, award – need an overhaul. There is no agreement as to what exactly needs to be done – there is quite a bit of research on the topic of research funding from around the world, with many different suggestions. This essay looks at the problem from what is perhaps a novel point of view – that funding projects is not the most effective way to fund research in basic science, or to produce better science – funding scientists may be a more productive approach. Broadly speaking, this statement has been made at least once before, as I found out while doing research for this essay, but the specific suggestions I make are quite different. The funding agencies I have in mind belong to governments – while there are private agencies which fund basic research, both the budget and the scope of funding must be very large to implement the idea in this essay.

Why do governments fund research?

Scientific research is a collective human endeavour – even the maverick scientist of popular mythos, unsocial and working alone, has to rely on theories and results discovered by their predecessors and competitors. Scientific knowledge, which is the fruit of scientific research, can thus be thought of as a product of collective global effort. Economists have argued that scientific knowledge is a durable public good – like free education (where it exists), free roads, and the rule of law – because it can be consumed by one person without reducing its value for another, and also because it is practically impossible to exclude anyone from enjoying its benefits.

Scientific knowledge not only drives our economy through its many applications, it also provides the foundation for technological innovation. Much of the knowledge which is used for these purposes falls in the category of *applied* science. Scientific research also produces the intangible goods of pure scientific knowledge, which may be useful only in satisfying human curiosity and may have technological relevance only as a spinoff. This kind of research is usually referred to as *basic* research. The boundary between basic and applied science can be blurry, as it is not unusual for basic research to find some application. The intention of the researcher is important however, so if the scientist is not actually searching for something that has immediate practical use, I will say that they are doing basic research, even when the result happens to be a technological innovation. We could also define research in basic science as driven by curiosity rather than towards a predetermined goal. The typical example is theoretical physics, almost all of which falls in the category of basic research, as does much of computational and experimental physics.

Research in science is one of the most economically efficient forms of government investment. The product of applied research often comes to us through private enterprise, which serves as the interface putting scientific knowledge to useful public service. For any government, it is in the national interest to stimulate the creation of private goods by funding scientific research. But it is a limited view which sees the utility of science only in terms of an economic balance sheet. Knowledge of particle interactions at very high energies or of rotation curves of galaxies millions of light years away do not directly contribute to national economies, but they are deeply alluring to the human intellect. Of course, the experiments which investigate physical phenomena require highly sensitive instruments, which in turn require cutting edge research in applied science and technology. But even apart from that, research to gain scientific knowledge for its own sake is a global public good, simply because of the stimulation it provides to intellectual creativity. It can also be a matter of great national pride to be ahead of other nations when it comes to basic research. Not surprisingly, all developed nations spend a significant fraction of their GDP on basic research in addition to applied science. Investment in technology research produces what we need in the immediate and foreseeable future, while investment in basic science is more of a long-term commitment.

Serendipity and goal-oriented research

The fruit of basic research can take a long time to ripen. Research in climate change was for many years perceived as pointless and even irrelevant to society, economy, and even to the rest of scientific disciplines. Some universities did persist with climate change scientists however, and their efforts have led to the discovery of processes in the Earth's atmosphere which create long-term threats to human civilization. The discovery of these processes was often serendipitous, as scientists tried to model the climate using mathematical equations and compared the predictions of those models with observational data. That was how the hockey stick graph of global temperatures was modelled – and extrapolated to the future – at the end of the millennium.

The popular notion of scientific research – focused and perseverant investigation of phenomena, often involving experiments – does not seem to allow for a role of luck. But the history of scientific research is replete with stories of serendipitous breakthroughs and startling revelations. The principle of buoyancy, among the earliest recorded principles of physics, is said to have struck Archimedes in his bath when he saw the water overflow as he entered the tub. The ring structure of Benzene was visualized by August Kekulé after he had a dream in which a snake was eating its own tail. Wilhelm Röntgen discovered X-rays quite accidentally while testing if cathode rays could pass through glass. Alexander Fleming discovered penicillin when he accidentally left a bacterial culture unprotected while going on vacation and on returning, found that an invading fungus had prevented the bacteria from growing. More recently, Harry Kroto chanced across the spectroscopic signatures of molecules found in interstellar space and recognized them as long chain carbon molecules. He got an astrophysicist friend to search for and find more such molecules in space and realized that they might have been created by carbon-rich stars. Then he convinced Richard Smalley and Robert Curl to blast graphite discs with laser to create long chain carbon molecules and found C_{60} – the buckminsterfullerene, resulting in a Nobel Prize for the three. Although X-ray and penicillin found widespread application immediately upon their discovery, they would qualify as basic research, as their discoverers were not looking specifically for things which could be immediately useful to society.

This is not to say that serendipitous discoveries do not occur in applied research, but they are harder to come by. Applied or technology research is almost always focused towards well-defined goals of creating something practical and useful, often before a preset deadline. This can channelize thoughts and efforts in specific directions and reduce the role of chance. On the other hand, much of basic research is also goal-oriented. For large collaborative experiments there is usually a broad goal, with the understanding that many scientific problems will be studied and many branches of science and technology may be involved. For example, the goal of the Large Hadron Collider (LHC) is the verification of different theories of particle physics, including the discovery and study of the Higgs particle and the search for supersymmetric partners of elementary particles. In the fifteen years since it was first started up, it has produced more than 15,000 research papers (counting only ATLAS, CMS, and ALICE, the three main research collaborations at LHC, according to the inspirehep.net website) in these and many other problems in high energy particle physics. Thousands of scientists work in these collaborations and many ideas are discussed among them to decide which ones will be tested and which will be discarded. Many of these scientists are employed elsewhere and must apply for research grants – to pay for their travel to CERN, as well as to pay assistantships to graduate students.

Smaller experiments in basic science – those conducted in labs run by usually a single scientist (Principal Investigator or PI) with the help of junior researchers – are also goal-oriented in general, for the very simple reason that equipment are purchased, and research assistants are paid, from research grants which have to specify the goals of the experiment quite clearly. Most scientific knowledge today come from such grant-driven research, although in many cases the equipment keep functioning and the experiments keep running even after the grant has expired. Much of theoretical research is also driven by grants – for research that requires a lot of computation, grants provide for the computers (or computer time) just like equipment for experiments. For pencil-and-paper theoretical research, the actual requirement for grants is less – they are needed for travel and sometimes for paying fellowship or research assistantship to graduate students, and sometimes for augmenting the salary of the PI. Even in basic research, most of the research done on grant money is done with specific goals in mind, since a mid-term evaluation will compare the progress of the project with the original goals. While curiosity plays an important part in any kind of research, there is not a lot of scope for venturing outside the boundaries set by the grant proposal. The promise of specific results made to the taxpayer to secure funding often leads to predictable research couched in hyperbole.

Is it possible to reduce the dependence of scientific research on grants, so that basic science is less strongly bound to fixed goals? Scientists can and do work on non-grant research problems – it turns out that 30% of all Nobel-winning science research between 2000 and 2008 were not funded by research grants. Reducing or denying research funding is clearly not a solution. Is there another option?

Funding competition

There is another, very important purpose of the research grant beyond the scientific value of the project – the ability to win grants is crucial for career advancement of scientists. Grants bring in rewards in terms of promotion, prestige, and professional accolades. Scientists are often judged by their institutions by the amount of money they bring in. Rejection of a grant proposal can cause mental stress and also means more time has to be spent in rewriting the proposal or in writing a new one. The time spent by scientists in writing and rewriting grant proposals was nearly a quarter of the time spent by them in doing research, in the top 150 research universities in the United States in 2004-5. In some fields, this fraction may even be higher. Assistant Professors spent 50% more time writing grants than full Professors, women scientists spent 30% more time than men. While it is true that writing a proposal can help a scientist focus thoughts and crystallize ideas, the time for writing is borrowed from the time to do research. It has been suggested that the total effort spent by all researchers in writing proposals may be comparable to the total scientific value of the research that is funded. There have been several suggestions of alternative funding procedures, but funding agencies have persisted with more or less the same model. In this essay I suggest yet another way to decide funding that reduces the time commitment for both the researcher and reviewer.

Research grants are awarded on the basis of competition in which reviewers evaluate and rank many grant proposals within a fixed time. The reviewers are experts in the field, usually senior people, which is both a good thing and a bad thing. They have the experience to determine if the project is doable, if the equipment, materials, and research assistants are sufficient (or in fact necessary) to achieve the promised goals, and also if those goals represent an advancement of science. On the other hand, senior scientists tend to be more fixed in their ideas about what needs to be investigated and by what methods, and what can be achieved in that. Homophily, or the tendency for people to be attracted to those who are similar to themselves, or do similar work to their own, creates another kind of bias in the peer-review process. As a consequence, research grants tend to accumulate at a small number of ideas. In the high energy theory and gravity landscape, some examples from recent history are supersymmetry, extended objects, loop quantum gravity, and similar topics – many of which have been extremely attractive and hugely productive in terms of papers and results, but so far have failed to connect to experimentally observed results. Outlying ideas are funded rarely, if at all. The reviewers also spend a significant amount of their time in evaluating projects, and their attempts to reduce this time can lead to unequal times being allocated to review different projects, and thus inequitable treatment of proposals.

It is widely recognized that the standard method of funding is not the most efficient way of distributing funds and several alternative methods have been suggested. One is a lottery system, which would take all applications satisfying some basic eligibility criteria and select by lottery the ones to fund. This removes the possibility of bias in the review process, but may not reach many who are doing interesting work and need the funding. Another is to evaluate applications using an automated procedure such as Machine Learning or Artificial Intelligence. This method avoids personal bias, but relies on bibliometric indices such as citations, impact factors, and the like, mined from databases. The problem with indices is that they are usually affected more by career statistics rather than recent work and can be gamed, while databases often have incomplete information, spelling errors, incorrect attribution to similar names, and many other shortcomings. A third option is to demand that project proposals have broad goals. This may make it easier to write proposals and to review them, while leaving open the possibility of serendipitous discoveries and innovations. On the other hand, broad goals invite exaggerated claims of what may be found, early career researchers may not be able to see beyond specific targets, and homophily would play a role in the evaluation more often.

None of these methods reduce the time commitment of the researcher. Except for the automated evaluation, the time taken for review is also not reduced greatly. In this essay, I will take the view that the most important purpose of funding reform is to free up the time to do research, as well as to loosen the commitment to specific goals.

A "no proposal" model

The format of the project proposal in basic sciences is fairly similar the world over. A scientific idea lies at the core of the proposal. The idea is embellished with a background – what makes the idea interesting, what has been done before that led to this idea, how it fits in with existing knowledge. It is developed as a project – what is planned to be done, what equipment and manpower might be needed, what kind of results it may produce, what kind of competition there is, nationally and internationally, to find such results. How the proposed PI is (ideally) suited to solve the problem proposed in the project, how their past work and that of their collaborators have prepared them to take on such a project. And finally the most important part – how much money will be needed and how it is planned to be spent – for purchasing equipment and materials; for travel by the PI and by their students and collaborators to perform experiments, for collaborative work, or to attend conferences; for sharing with the parent institution for overhead expenses; for paying fellowships to students and wages to research assistants; for augmenting the salary of the PI if necessary, etc.

Most of the time spent in writing a proposal goes in the scientific portion, as should be expected. The proposal has to convince the reviewers that the project is grounded in existing scientific knowledge, is likely produce useful results (or at least interesting ones), and that the research background of the PI makes them eminently qualified to work on the project. Most of the time needed by reviewers also goes into these portions – regrettably, these are also places where some bias can creep in. For example, the reviewer may feel slighted if their work was not adequately cited – this can affect their evaluation of the project. The reviewer may be particularly interested in supporting some specific lines of research because of homophily, or not willing to support some research because of a lack of personal interest. Or the reviewer may not have a lot of faith in the PI as a researcher because of various hidden biases, which may or may not have anything to do with the science being proposed. It is also known from actual experiments that if the team of reviewers is changed, the grades or ranks received by proposals change completely. While there is no foolproof way of removing all bias where human evaluation is involved, it is possible to smoothen out most of it by having many reviewers look at every proposal and averaging the grades given by them. But the time required to review each proposal makes this impossible.

The "no proposal" model proposes to radically simplify the grant application-evaluation-funding chain in basic science research, by removing most of the scientific part of the proposal. In this model, the application for a 3-year grant will include only the field of research, the list (and perhaps abstracts) of publications of the PI over three years and an outline of a budget. The most important part of a grant, from the perspective of the PI, is the money allocated, as well as any restrictions on how that money can be spent. It is then expected that the PI can always provide a reasonably clear budget outline – specifically mentioning equipment which need to be purchased, details of travel expected in three years, number of students or research assistants needed in that time and what they will be paid, overhead costs for the parent institution, etc. But nowhere will the PI need to specify what science will be done with the funding, leaving themselves free to investigate whichever problem takes their fancy. Perhaps they will continue to work on things similar to what they have done before, or perhaps they will find explore new avenues of research in parallel, without being in great danger of losing funds. It should be easier for researchers to start interdisciplinary research or to shift disciplines in this way. It will also be easier for early career researchers to secure funding, as only recent performances will be considered and compared.

The reviewers should also find it easier to grade the applications. What they will need to consider is the performance of the PI in the past three years and from that make a reasonable guess about their performance in the next three years. This can be a vexatious problem, but it is also partially solved, since each publication has already gone through one or more referees, and may have been cited by others, so the evaluation is to some extent crowd-sourced. The funding agency will likely prescribe how to grade recent performance depending on what kind of research it wants to fund – it could include some weightage of the number of publications, number of citations, impact factors of the journals, interdisciplinarity, and also subjective opinions such as broad interest, current relevance, and novelty. Guesses of the future performance may be framed according to similar parameters as well. But reviewers will not know exactly what problems the PI is thinking of solving. So they will have to have faith – that the applicants they have backed will produce *something* of value to science which are not lesser in quality than what the present format produces.

Human evaluations can never completely avoid charges of bias or favouritism, while machine evaluations suffer from various shortcomings as we know. But since there is no proposal to be judged, each application can be graded fairly quickly after first separating out according to the field of research and the money requested. A large pool of reviewers can be asked to grade all applications in each category, hopefully washing out fluctuations due to bias. There are additional ways of minimizing bias that are currently known but not used very often. For example, a lottery can be used after a broad ranking in each category. Machine evaluation can be used more efficiently as well. Funding agencies usually have a system of mid-term evaluation, which compares achievements against expectations. Once the model is running for a few years, the data could be used to set up a Bayesian model to help guess future performance.

Funding agencies spend billions of dollars in funding basic research, but there is a general consensus that the funding process is far from efficient. The point of the "no proposal" model presented here is to allow researchers to focus on their science as well as to work on what they please and not be held to a commitment made primarily to secure funding. Admittedly it sounds very odd if funding agencies report that a thousand scientists have received research grants, but it is not clear what research they will do with it. It is much more inspiring to hear about a thousand avenues of research which have been funded. One possible way of thwarting criticism on this ground is to think in terms of prizes rather than grants – or "research prize" instead of "research grant." It sounds more prestigious, the word "prize" indicates that it is a reward for past work, and the word "research" hints that it is meant for research and related expenses. However, it also reduces the control exerted on research by funding agencies. For that reason, the "no proposal" model may not be easy to implement. But if this model is adopted by any agency, that will surely change the way science is done, for the better.