ISRAEL JOURNAL OF ECOLOGY & EVOLUTION, Vol. 56, 2010, pp. 1–7
DOI: 10.1560/IJEE.56.1.1

IJEE Soapbox: COOPERATION, COMPETITION, AND THE SOCIAL ORGANIZATION OF THE SCIENTIFIC ENTERPRISE

Robert D. Holt
Department of Biology
University of Florida, Gainesville, Florida 32611-8525, USA

Winston Churchill once famously quipped, "It has been said that democracy is the worst form of government, except all the others that have been tried." I sometimes wonder if the same holds for how public funds are used to sustain the scientific disciplines of ecology and evolution, and more broadly the scientific enterprise as whole, through agencies such as the National Science Foundation (NSF) in the United States (US) and the Natural Environment Research Council in the United Kingdom. In this essay, I will put down some reflections that have occurred to me over time about this pragmatic side to the study of ecology and evolution, and also about how some aspects of the social organization of scientific funding implicitly reflect the kinds of mechanisms which ecologists, evolutionists, and behaviorists study in their own domains.

In an open society, our scientific knowledge of the world, taken collectively, is what economists call a "public good" (Gravelle and Rees, 2004; Stiglitz, 1999). A public good is something, such as a weather forecast on the radio, which cannot be made unavailable to one person because another person "consumes" it, and over which individuals cannot exert "ownership". In other words, in ecological jargon, there should be neither exploitative nor interference competition over public goods.

So my learning about, let's say, the potential for chaotic dynamics in even simple nonlinear ecological models, or stable isotope techniques, or how to tell a braconid from an ichneumonid, does not necessarily make it harder for you to learn something about the same topic, too. Indeed, if anything, there is potentially a kind of mutualism in learning, in that the larger the community of individuals interested in a particular subject, the more likely there will be an inspired author or two crafting textbooks and popular treatises that can lay out the lineaments of that subject in a lucid manner. Students learn from

---

E-mail: rdholt@ufl.edu

The IJEE Soapbox provides an informal forum for leading ecologists and evolutionary biologists to expound on issues that they find particularly exciting, thought provoking, and novel.

Robert D. Holt is our first invited ***IJEE Soapbox*** essayist. Bob is Professor of Biology and Arthur R. Marshall, Jr., Chair in Ecology at the University of Florida, and is one of the foremost theoreticians in ecology and evolutionary biology. His research focuses on theoretical and conceptual issues at the population and community levels of ecological organization and on linking ecology with evolutionary biology. Bob is best known for his pioneering work on apparent competition, multispecies interactions in food webs (community modules) in time and space, and the evolution of niche conservatism.

each other as much as or more than from their instructors and texts, so having compadres struggling with the same material can help cut through the clouds of confusion that are almost always present in difficult areas of science. In some circumstances, of course, learning requires specialized equipment, first-hand access to a gifted teacher, or slots in a restricted course, so that an element of competition for limited resources can creep into the learning enterprise itself. This competitive element is automatically present for courses graded on a curve or where professional certification mandates competition for a limited resource (e.g., for admission into medical schools). Some kinds of knowledge come only from practice in a kind of apprenticeship relationship, akin to learning an art form, and so it is automatically less true that such knowledge is a public good. A rule-of-thumb might be that the kind of knowledge, including scientific knowledge, that is a public good is the sort that can be written down in a published paper, monograph, or textbook, or otherwise recorded (e.g., audio and video, or web-based) but there are other kinds of knowledge that are experiential and personal, and so perforce private goods. There is apprenticeship and artistry, even in science.

Understanding how competition and cooperation operate across a hierarchy of levels of organization is fundamental to understanding the history of life (Maynard Smith and Szathmary, 1995). Morris (2009) has suggested that selection among individuals within academia contending for faculty positions should be viewed as a multi-level selection process. Likewise, the interplay of competition and cooperation pervades the social and economic organization of science, in terms of how funds get allocated among investigators and projects. Ironically, the painstaking bit-by-bit building of the corpus of scientific knowledge (a quintessential public good, taken as a whole) typically involves fierce competition for "private goods"—in particular, funding for research projects. If I eat a sirloin steak, you cannot eat that very same steak; the steak is thus a private good. And if Dr. Q through her university gets a grant from a funding agency, those exact dollars or shekels or euros will not be going to Dr. P at his university. The money comes from the public, but nonetheless, grants are in effect private goods.

Societies, and indeed subsocieties within societies, differ greatly in how it is determined who gets which private goods. In some cases the mechanism is little more than despotism, or more benignly, nepotism (including inheritance). In idealized free markets, ownership of goods is determined by some kind of exchange that is freely entered into, rather than coerced, between the participants. In the case of NSF and other like governmental funding agencies, there is often an interesting blend of cooperation and competition. Competitive mechanisms for funding research certainly are the base of the scientific enterprise, in much of the world. In the US, the predominant agency that supports research in the basic sciences of ecology and evolution is NSF, and so my thoughts will be focused largely on it. For the most part, the only way a scientist (via his/her institution) extracts research funds from NSF is to write a proposal, which is then submitted to a particular program, where it is evaluated by a panel of peers in terms of its relative merit. The final resolution of funding is then made by a program officer (different agencies have somewhat different procedures). There is, in effect, a near-Darwinian competitive struggle among grant proposals arriving before any particular panel, competing for a relatively modest pie of resources. Only a small subset will remain standing at the end of each round of competition.

But the actual mechanism of determining these “winners” involves a substantial amount of cooperation within the scientific community, via the reviews and perspectives provided by the panel itself, and external reviews. These individuals are not at all handsomely rewarded for their participation as adjudicators in this process, at least not in any direct sense. Instead, this social arrangement is a blend of what appears to be pure altruism (e.g., providing reviews for a foreign funding agency, where I know neither the scientists nor the program officers), enlightened self-interest (e.g., one learns from the process), and a form of “indirect reciprocity” (Nowak and Sigmund, 2005). By doing an honorable and fair job as a reviewer or panel member, my expectation is that others will do the same for my own proposal when it is submitted (see McPeek et al., 2009, for a similar point about the reciprocal altruism that is required for the scholarly peer review system to function). For such indirect reciprocity to work depends on reliable social acquaintance and reputation; in the case of NSF and other agencies, this indirect reciprocity is typically mediated through the knowledge program officers have about their own scientific communities.

Having served over the years on many panels, and submitted a fair number of proposals myself, I can attest to the fact that taken as a whole, the process is reasonably fair, within the overall constraints of there being far too many proposals submitted to possibly be funded in any given round. I have occasionally been surprised by how panelists from rather disparate backgrounds can converge on relatively comparable opinions of given proposals during panel discussions (but see below). Renowned scientists can be shot down, and excellent proposals from young scientists have at least a modest chance. Program officers in general try to be fair-minded, and carry a large work-load with rather little staff. The act of writing a proposal in itself can help a scientist clarify his or her ideas. And so on. A lot of terrific science has certainly been supported by our granting agencies. Maybe the system is about as good as could be expected.

Having made these points, all of which are to the good and speak well of our current funding system in science, I think one could craft a counter-argument that in the current system there is also a great deal of wastage of human time and effort. If the corpus of scientific knowledge as a whole is a public good, then we should ask how the social organization of science and in particular how funds are allocated can be optimized, relative to some “utility function”. I do not mean this in a literal sense, but at least figuratively. There are many directions this argument could go.

The primary urge for writing proposals is or should be an intellectual objective, but increasingly, administrators use the amount of grant funds generated by faculty to determine hiring, tenure, and promotion, and even annual raises and space allocations. These external pressures can lead to a kind of pressure cooker atmosphere among all faculty, but junior faculty in particular. As Marshall et al. (2009) note, successfully landing grants is becoming essential even for entry-level academic jobs in ecology and evolutionary biology. I think we have a poor handle as yet on the distortions that have arisen in the academy because of the perception by administrators that scientists can be "cash cows" for their institutions.

Let us consider the potential life histories of a proposal submitted, say to NSF. There are three possible fates for any proposal which has been submitted for the first time: i. It can be funded right off the bat (Yes, Virginia, this does happen); ii. It can be funded, but only after one to numerous resubmissions; iii. It is never funded, even with resubmission.

At present, only around 10% of the proposals that go to the Population and Community Ecology panel at NSF (as one example) get funded in any given round. As was said in another context: "Many are called, but few are chosen" (Matthew 22:14). Since this includes resubmissions, funding upon first submission is thus surely a quite low percentage. Out of the 90% or so of proposals that do not get funded in any given round, there are surely a few which have fundamental flaws such that the expenditure of public funds would be a waste, but in my experience the great majority are actually quite fine science. Let us say that a modest percentage, around 50%, actually are good science. Then for all these proposals to be funded eventually would require five iterations of that panel, assuming no new proposals came in. (But then of course there are always new proposals, with which resubmissions have to compete.)

Some of the proposals which are denied in one round continue to be polished and improved in subsequent submissions, until finally at long last the proposal is funded, and everyone involved—the investigator, the panelists and reviewers, and program officers—gives out a big sigh of relief. Voilà, the process is vindicated! But wait a minute—is the science actually all that much better when it is finally carried out, than it would have been earlier on? We would all like to think so. Yet the definitive study, to my knowledge, has never been carried out (it is rather hard to conduct the appropriate controlled experiments!). By analogy, students in the US may re-take the SAT (a standardized test required for admission to many colleges and universities), and in so doing may improve their scores, and many students take specialized courses which presumably prepare them for the test. But the efficacy of such measures may reflect not a deepened understanding of English, math, and so forth, but simply a sharpened set of skills for taking tests per se.

"Grantsmanship" denotes the art of obtaining grants, which involves a skill set which may or may not closely match the talents and drive required to carry out scientific work at a conceptually deep and innovative level. I suspect any active scientist can think of some individuals who are very successful over their careers at landing grants, yet for whom it is hard to point to substantial scientific advances resulting from those grants, and to other individuals who have made deep and enduring contributions to science, with a much spottier record of grant success. Highly creative ideas can come across as "risky" research, and often on first public showing will have an assortment of warts not yet removed.

Any time spent writing a proposal is perforce time not spent carrying out experiments or simulations, or writing papers, or advising students, or even just thinking deeply about one's subject beyond the immediate focus of the proposed research, so there is a real missed opportunity cost in preparing a proposal. I have heard colleagues estimate that some hundreds of person-hours are required per NSF proposal. The more resubmissions are required before a project is funded, the larger is this opportunity cost. Long-term studies often are particularly difficult to keep afloat, because they require recurrent bouts of funding, and there is a risk it will not materialize on schedule; a colleague recently told me about one such study (which has had a sterling record of publication success), which required ten (10!) submissions in various guises to keep it alive. Likewise, if a project to be successfully executed requires a particular constellation of investigators, even a single failure in landing a grant makes it harder to keep those teams assembled, because of the centrifugal forces acting on people's lives. The process itself thus arguably selects in subtle ways against certain kinds of important research programs.

And proposals for lines of work which are in the end abandoned, after one to several attempts, are almost a complete waste of time, aside from being dispiriting. One of the real challenges is that science, as with any human endeavor, has its fads and fashions. Research agendas that seem out-of-fashion at the moment are bound to run up against brick walls in panels, but may nonetheless be very valuable in the long run. Yet they have a very hard time being funded because of this element of social selection in favoring certain proposals over others. This aspect of the process is a bit reminiscent of sexual selection, which can lead to outcomes in evolution that in a broad sense are maladaptive for populations. One of my scientific friends jocularly refers to landing a grant as winning a "beauty contest". That may be a bit strong, but there is an uncomfortable element of truth to the jibe.

I think there is a range of research questions here for some astute social scientist, to assess the magnitude and societal implications of these opportunity costs, in terms of wasted human hours and potential, borne at the level of individuals and filtering up to the scientific enterprise as a whole. There should also be research on how the social dimensions of science influence the selection of proposals.

Another question that might be addressed by such a social scientist is the degree to which the adjudication of which proposals are funded, versus which are not, reflects "deterministic" factors that in some sense are objective, versus "stochastic" factors. How dependent is the final outcome on the exact composition of the panel, and of the reviews that happen to have been submitted, and other chance factors, versus intrinsic and objective qualities of the proposal and investigator? It is easy to slide into a Whig version of history (Butterfield 1965), where the final "winners" are viewed as in some absolute sense necessarily better than the "losers". But a little intellectual humility is, I think, in order; there are error bars around all human decisions. Consider the following thought experiment. Imagine that there is a pool of 100 potential members of an ecology panel, and that the normal panel membership is 20. If we were to carry out a controlled experiment, we would try to establish a number of randomly formed panels of this size from the original set (allowing for stratified sampling to represent subdisciplines, gender diversity, etc.; the sampling would need to be without replacement, so as not to totally exhaust certain panelists!). Then, give them all the same set of proposals to review and rank independently, carry out a rank–order correlation for each pair of panels, and then average these scores. I expect there would be a reasonable amount of commonality among the rankings, but how much? It is unlikely to be either 1 (complete congruence), or 0 (complete unpredictability across panels), but what if it is, say, 0.6? In this case, stochastic elements would appear to play a very large role in the short-term course of science, as governed by the allocation of funding, since many proposals would not have been funded in some panels, yet would have been funded in others.

An experiment of this sort was actually performed in the 1970s (Cole et al., 1981), not in ecology, but in chemical dynamics, economics, and solid-state physics. These authors concluded that "the fate of a particular grant application is roughly half determined by the characteristics of the proposal and the principal investigator, and about half by apparently random elements which might be characterized as the "luck of the reviewer draw". There is no particular reason to think that the result would be any different in ecology and evolution.

Understanding the relative importance of stochastic versus deterministic processes is, of course, a fundamental issue across the ecological and evolutionary sciences, from understanding the determinants of community assembly, to the generation of genetic differences across clades. I suggest that it may be instructive to critically view the social organization of science and how funds are allocated through the same lenses. There is bound to be considerable stochastic fuzz around a deterministic signal.

When I asked a colleague at NSF (who wished to speak off-the-record) about the nature of the "utility function" going along with the presumed public good of scientific knowledge that the agency promotes, he jokingly replied that "A utility function assumes that the system exists as it does for a rational reason and that there is some sort of feedback mechanism between outcomes and the shape of the function. I highly doubt either :-)."

These are good points. As for any institution, there is a historical dimension to how NSF functions today that still reflects the particular social, economic, and indeed world conditions at the time it was created. The fact that the system promotes the funding of projects, rather than, say, persons (which is more the Canadian model), may be so hardwired that it is pointless for us to think much about how it might be differently arranged. There are disadvantages in funding persons vs. projects. First of all, "the rich get richer". Chance vicissitudes early in a career have a way of magnifying themselves, throughout. By analogy, if there is an exponential growth process without density dependence (in population biology, or bank accounts), with the growth rate a stochastic variable, the ultimate distribution of outcomes (population size, or capital) has a lognormal distribution, with a long tail of exceptionally "wealthy" individuals. In like manner, scientists who have a "run" of good years, in terms of publications and research findings, will necessarily look better than peers, who simply were not so lucky. This then sets them up to compete more strongly for grants. Secondly, what a person does in his or her career depends not on just who they are, sui generis, but as much on where they are, and who their associates are. Any individual is a node in a social web, and how strong and creative they appear to be to the outside depends in part upon the structure and strength of the web in which they are embedded.

In recognizing the importance of historical constraints in the social organization of science, it might be analogous to the human backbone. When our ancestors became bipedal, they inherited a vertebral structure that made many of us prone to chronic and often very painful back conditions. We cannot really do anything radical about this, but instead just cope (as for myself, I get physical therapy, massages, and regular exercise). So with respect to the issue I raise above, about opportunity costs and so forth in our present system of funding research, maybe the best we can hope for is a form of modest meliorism (the philosophical stance that one can with effort make progress in the world, ameliorating conditions, without aiming at creating Utopia)—tinkering with the system to smooth over and minimize such costs and biases, trying to improve matters without pretending to overthrow anything in a basic way.

This seems like a good place for this essay to stop. Anyway, I need to go work on my next proposal, which I am sure will be quite a pain in the—backbone.

ACKNOWLEDGMENTS

I thank Sam Scheiner, Alan Tessler, and Doug Levey for congenial interchanges on the topic of this essay, Larry Kenny for insights into the usage of terms from economics, Mike Barfield, Leon Blaustein, and Burt Kotler for thoughtful remarks on the text, and the University of Florida Foundation for its support.

REFERENCES:

Butterfield, H. 1965. The Whig interpretation of history. W.W. Norton, New York, NY.

Cole, S., Cole, J.R., Simon, G.A. 1981. Chance and consensus in peer review. Science 214: 881–886.

Gravelle, H., Rees, R. 2004. Microeconomics, 3rd Edition. Prentice Hall, New York, NY.

Marshall, J.C., Buttars, P., Callahan, T., Dennehy, J.J., Harris, D.J., Lunt, B., Mika, M., Shupe, R. 2009. In the academic job market, will you be competitive? A case study in ecology and evolutionary biology. Isr. J. Ecol. Evol. 55: 381–392.

Maynard Smith, J., Szathmary, E. 1995. The major transitions in evolution. Oxford University Press, Oxford, UK.

McPeek, M.A., DeAngelis, D.A., Shaw, R.G., Moore, A.J., Rausher, M.D., Strong, D.R., Ellison, A.M., Barrett, L., Rieseberg, L., Breed, M.D., Sullivan, J., Osenberg, C.W., Holyoak, M., Elgar, M.A. 2009. The golden rule of reviewing. Am. Nat. 173: E155–E158.

Morris, D.W. 2009. Life, history, and multi-level selection in academe. Isr. J. of Ecol. and Evol. 55: 393–394.

Nowak, M.A., Sigmund, K. 2005. Evolution of indirect reciprocity. Nature 437: 1291–1298.

Stiglitz, J.E. 1999. Knowledge as a global public good. In: Kaul, I. Grunberg, I., Stern, M.A. eds. Global public goods: International cooperation in the 21st Century. United Nations Development Programme, Oxford University Press, Oxford, UK, pp. 308–325.