Modelling academic research funding as a resource allocation problem

Nicholas Geard and Jason Noble

School of Electronics and Computer Science, University of Southampton, UK
[nlg|jn2]@ecs.soton.ac.uk

Abstract. Academic research funding is allocated through a competitive bidding process that may lead to inefficiency as excessive time is spent on proposal writing. We develop a simple agent-based model of the process and find that current systems are indeed likely to be inefficient. Alternative allocation schemes involving either a cap on individual effort or appraisal from the centre are indicated as improvements.

1 Introduction

Many organizations and institutions face resource allocation problems [7]. Companies would like to ensure that the greatest expenditure goes to those divisions best able to make a profit, the military want new equipment to go to units best able to make use of it, and schools and universities want scholarships to go to the brightest applicants. Some allocation problems are simple because identification of the best candidates is straightforward. In other cases, there can be considerable administrative or cognitive effort involved in determining just who the best candidates are. The agent or agency responsible for making the allocation (*e.g.*, the central administration of a company) may naturally be tempted to off-load some of this effort onto the agents that compete for the allocated resources. This is often achieved by introducing a bidding system. For example, in a corporate context, by asking each division of a company to set out a case for their resource needs, and using the strength of these written cases to determine the appropriate allocation. This can be contrasted with the central administration investigating for themselves the relative efficiency and profitability of each division and determining the allocation that way.

The introduction of a competitive bidding system means that candidates for the receipt of centrally allocated resources need to decide how much time and energy to put into their normal activities, and how much into the production of bids. The strategic situation starts to look like the costly signalling games that evolutionary biologists have identified in sexual selection contexts [8, 15]. In these games, a female has a valuable resource to allocate: a mating opportunity. Males engage in competitive signalling that is designed to convince the female that they are the best candidate. Examples include the elaborate tail growth and associated display behaviour of peacocks, and energetically costly croaking by male frogs. Game-theoretic analyses have shown that if the marginal costs of a stronger signal are higher for the less able males, then honesty will be enforced. It will not be worthwhile for low-quality males to give the same level of signal that the best candidates can give, and thus the female can reliably mate with the best males by attending to the quality of their signals.

Anecdotally, competitive bidding in social and organizational contexts can be quite similar. The production of a competitive bid is an expensive process that reduces the time candidates can spend on their primary functions. The time "wasted" in this activity may appear to be a necessary evil if it means that the best candidates are successfully identified and awarded the resources. Note, however, that although the costly signalling systems found in nature are evolutionarily stable, this does not mean that they are efficient. If we are dealing with a designed system rather than a naturally evolved one, the time and energy involved in working up a competitive bid or sexual display could be dispensed with if there were some way of enforcing honesty through a cost-free signalling channel. For example, if candidates somehow wore their quality on their sleeve in a way that could not be faked, those allocating the resources would be able to attend directly to this information and costly bid production would be unnecessary. Overall system efficiency, measured in terms of the time and energy devoted to primary production as opposed to bid preparation, would be increased.

The current paper focuses on the particular resource allocation problem associated with the funding of academic research [12, 10]. We will focus primarily on the situation in the UK simply because we are more familiar with it, but similar systems exist in other countries. In the UK, seven research councils distribute approximately three billion pounds per year in research funding [3]. All of this funding is allocated through a competitive bidding process: academics write grant proposals and send them to the appropriate research council, the proposals are reviewed for quality, and a panel of experts decides which proposals will be funded. Some basic research capacity is supported by core funding that is directly allocated to each university, but the proportion of UK research output that is paid for in this way has been decreasing for many years [13].

The competitive bidding process for research funding means that a significant proportion of a researcher's time is dedicated to applying for funding to carry out research, rather than to research itself [11]. Increasingly, this is true not only for resource-intensive research areas such as biomedical sciences and high-energy physics, but also for the less resource-intensive subjects such as social sciences and the humanities [4, 2, 9]. Most research proposals will fail to secure funding for their authors: for example, the UK's Economic and Social Research Council reported that only 19% of proposals were funded in the academic year 2008/2009, and the Engineering and Physical Sciences Research Council funded 26% of proposals over the same period. The time and energy dedicated to the production of the failed proposals in these cases has been wasted.

Individual academics have to make strategic decisions about how much to time to spend on their primary research versus preparing funding proposals, and in some ways they are in a similar strategic predicament to peacocks and frogs. If all academics spent some modest proportion of their time (say 5%) on bid preparation, the best proposals would be funded and all would be well. However, this would potentially open up a strategic advantage for those prepared to spend 10% of their time on preparing bids, and the route to an inflationary spiral of effort becomes clear. In some contexts academics are indeed spending a surprisingly large proportion of their time on preparing grant proposals. Schaffer [14] reports that "At the University of Pennsylvania School of Medicine, for instance, faculty members normally spend about 50 percent of their time working on grants... in March and April, however, faculty members have spent more like 75 percent to 90 percent of their time going after stimulus dollars." Statistics such as these raise obvious questions as to whether this use of researchers' time and energy is optimal, and whether a more efficient approach to resource allocation could be pursued in this sector: for example, Gourdin and Poulin [5] estimate that skipping the bidding process and distributing equal funding to all eligible researchers would be more efficient.

In this paper we develop a simple agent-based model of the academic research funding system. Our aims are to investigate the level of inefficiency in the system assuming that each academic acts in their own interest, and to use the model to explore possible alternatives to the current competitive-bidding approach. Our intention is not to provide a precise match to how any particular funding system behaves now, or would behave after a proposed change; in the terminology of Heath [6] we are fulfilling a "mediator" rather than a "predictor" role. Although a predictive model would certainly be helpful, we agree with Agar [1] that a premature insistence on quantitative precision in models of social processes can be misplaced.

2 The simulation model

We model a population of academics ($N = 100$) competing for funding. As in the real world, funding is limited: there are only enough grants to fund a fraction P of the population; here, we use 40%. These academic agents are not all alike: each possesses an underlying research talent level R that determines their output (*e.g.*, expected number of journal publications per semester). In the current experiments, R ranges between 0.0 and 1.0, and is uniformly distributed.

Research funding is distributed once per semester. Grants are awarded to individuals, who can only hold one grant at a time. Being awarded a grant increases an individual's research output by a factor G; here, we use 125%. We assume that more talented agents are able to make better use of the extra resources. So, for example, an agent with $R = 0.5$ would produce 0.5 units of output per semester with no extra funding. If the same agent was awarded a grant, output would increase to $0.5 + 125\% \times 0.5 = 1.125$ units. However, in a competitive bidding system, agents must allocate some proportion A of their time to writing proposals, which reduces the time they can spend on personally doing research. In general, an agent's research output O is given by $O = R(1-A)$ for an agent not holding a grant and $O = R[(1 - A) + G]$ for an agent holding a grant. An agent's strategy is then the proportion of its time that it spends preparing applications as opposed to doing research (*i.e.*, A); this can range from 0–100%, in 10% units. Agents at the start of a run are initialized with random values of A between 0 and 50% inclusive.

Each semester, all agents with $A > 0$ submit a grant application, which is then evaluated by the funding body. We make four assumptions about this evaluation: First, that the perceived quality Q of a grant application in some way reflects the research talent R of the applying agent (if it did not, there would be no reason for the funding body to pay attention to applications). Second, that an agent with less research ability can compensate by investing more time in writing a persuasive application. Third, that there are diminishing returns on time invested (implemented here using the nonlinear *tanh* function). Finally, as no decision procedure will be perfect, we adjust the perceived application quality by a modest noise factor. Application quality Q is therefore given by $Q = [R + tanh(2A)] \times (1 + \eta)$ where η is represents Gaussian noise with mean 0 and standard deviation 0.1. Submitted applications are ranked by perceived quality Q, and the authors of the NP highest quality applications are awarded grants. Each academic then produces a quantity of research as described above.

Between semesters, agents adjust their strategy, reflecting the strategic tuning of proposal-writing effort by real academics, each keen to publish more papers than their rivals. Perhaps unsurprisingly, this was a subject on which we could find only anecdotal data, and so we defined and compared several heuristic strategies that capture what we feel to be key variables.

THERMOSTAT: Our simplest agents update their strategy much as a thermostat adjusts energy used in heating. Agents who obtain a grant in the previous semester reduce the amount of time they allocate to application writing by one unit (*i.e.*, "I succeeded in getting a grant last time, so perhaps I can get one again with less effort, and thus spend more time on direct research"). In contrast, agents who fail to obtain a grant increase the amount of time they allocate to application writing by one unit.

MEMORY: Exploratory runs indicated that systems using the THERMOSTAT heuristic exhibited several implausible behaviours: successful agents would oscillate between obtaining a grant and failing to do so as they honed in on the optimal amount of time to allocate to proposal writing. Untalented agents who never obtained a grant would continue to allocate 100% of their time to writing applications with no hope of ever succeeding or reducing their efforts. Both of these behaviours, particularly the second, could be seen as irrational. Our goal was to model reasonable behaviour by the academic agents so that we could ask questions about efficiency in the overall system. We therefore investigated a more sophisticated agent that had a memory of its past performance and operated according to the following rules:

1. If I get a grant, but can recall not getting one during the recent past, keep my proposal-writing efforts at the same level (*i.e.*, "don't get cocky").
2. If I have consistently got grants during the recent past, I may be putting in more effort than I need to, so decrease my efforts by one unit.

3. If I fail to get a grant, but can recall getting one in the recent past, I must be somewhat competitive, so increase my effort by one unit.
4. If I have consistently failed to get grants during the recent past, then drop out. From now on, submit no proposals and focus on my own research (with some small probability of re-entering the system in each future semester, with $0 \leq A \leq 0.2$).

Variants on the MEMORY heuristic can be characterized by the size W of the memory window defining "the recent past" and the probability E per semester of a dropout re-entering the grant-writing contest. We looked at two variants: MEMORY A: $\{W = 5; E = 0.05\}$; and MEMORY B: $\{W = 3; E = 0.02\}$.

FIXED: As an alternative, we investigated a scenario in which agents do not have an individual strategy; rather each agent spends a fixed 10% of their time applying for grants. An alternative view of this could be that the funding body performs some independent evaluative function (*e.g.*, bibliometric assessment) that results in an administrative cost equivalent to 10% of every agent's time. Either way it is impossible for agents to "cheat" by investing more than the allowed 10% of their time budget.

Evaluating system performance: We define a measure of system performance, return on investment (ROI), as follows: we calculate the summed output of the population, averaged over the final 10% of semesters in the run; this is the total output. We then calculate the amount of research per funding period that could have been obtained by the agents in the absence of any research funding (*i.e.*, assuming 100% of time spent on research, and no grant bonuses); this is the base output. The ROI is the difference between these two figures and it quantifies how much research the funding council's money has produced, over and above the research the academics would have done anyway.

3 Results

Analysis of bounds: Before simulating, we calculated several theoretical bounds on total research output:

$$O_{total} = \sum_{i}^{NP} R_i[(1-A)+G] + \sum_{j}^{N(1-P)} R_j(1-A) \qquad (1)$$

One baseline is the situation in which no funding is made available (*i.e.*, $P = 0$ and thus $A = 0$ for all agents). In this case, the total research output resulting from the agents' unaided research efforts is 50 (assuming a perfectly uniform distribution of R). A second baseline is the situation in which funding is allocated at random, in which case expected total research output is 75 (ROI = 25). We also consider the hypothetical scenario in which the funding agency is allowed perfect knowledge of each agent's research talent at no time cost to the agent. In this ideal scenario, total research output is maximized: grants are allocated to the top 40% of the population, nobody spends any time writing proposals, and the total research output is 90 (ROI = 40).

Next, we consider a somewhat more realistic scenario in which each agent spends an equal amount of time applying for grants in order to signal their research aptitude to the funding agency. As time spent applying for grants increases, the results for best and random allocations decrease linearly. If competitive allocation is to provide any benefit, it must lead to higher total output than random allocation. This is true only if agents spend no more than 30% of their time applying for grants (under this set of parameter values). This assumes that proposals transmit a perfect signal of underlying research talent to the funding agency: if grant evaluation is noisy the maximum time limit such that competitive bidding provides any benefit will be even lower.

Simulation results: Each of the four strategy variants were run for 100 simulated semesters; this proved long enough for the system to stabilize. Numerical results were averaged across 10 experimental runs with different random seeds, taking data from the final 10 semesters of each run (Table 1). The allocation methods were all generally successful in the narrow sense that grants were awarded to the more talented agents on average (*i.e.*, there was a positive correlation between R and success).

Table 1: Results for four different agent update strategies, showing means and standard deviations over 10 experimental runs.

	ROI	Mean O	Mean A	Corr(R, success)
THERMOSTAT	11.46 (1.62)	60.06 (3.07)	0.55 (0.02)	0.28 (0.02)
MEMORY A	13.53 (2.14)	63.74 (1.76)	0.37 (0.03)	0.68 (0.02)
MEMORY B	26.31 (1.56)	75.85 (3.40)	0.14 (0.02)	0.54 (0.03)
FIXED	34.76 (1.08)	84.97 (3.38)	0.10 (0.00)	0.83 (0.03)

The THERMOSTAT heuristic led to spiralling competition in the time invested in proposal writing, with the average agent spending 55% of its time applying for grants. Grants do tend to go to the more talented applicants, but there is a strong negative correlation between research talent and time invested on proposals: this was caused by the consistently unsuccessful low-talent agents putting in 100% effort. ROI was very poor at 11.46 units; significantly lower than random allocation (Figure 1), indicating that the excessive levels of proposal-writing effort here are truly pathological.

By allowing unsuccessful agents to drop out of the contest, the MEMORY heuristic limited the escalation of time invested in grant proposals. The A and B variants are superficially very similar, but Table 1 and Figure 1 show that they led to quite different results. MEMORY A performed little better than the THERMOSTAT heuristic, with a mean time allocation of 37% on proposal writing. MEMORY B looked much healthier, with ROI slightly above random allocation, and only 14% of the average agent's time being spent on proposal writing. MEMORY B had a shorter memory window and a lower chance of re-entry after

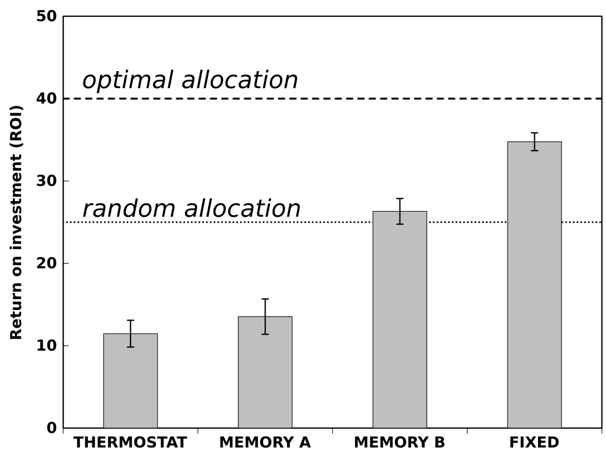

Fig. 1: Plot of mean ROI for the different heuristics. Error bars show standard deviations across 10 runs. All competitive bidding conditions had ROI scores worse than or at best comparable with random allocation of funding

dropping out, and thus the difference is explained by noting that the more time untalented individuals spend out of the system, the less pressure exists on remaining applicants. In the extreme case, so many agents drop out that average acceptance rate for a proposal rises to around 90%.

Results for the legislated time limit on proposal writing of 10% show a much improved picture. The correlation between talent and success is very strong at 0.83 and the ROI score of 34.76 units is much higher than any other condition. It is also a significant improvement on random allocation.

4 Discussion

The competitive bidding systems used to allocate academic research funding allocation inevitably lead to many hours of apparently wasted time when the majority of proposals are unsuccessful. Is this waste a necessary side-effect of a tough but fair process that sees the most money go to the best people? Or is it genuine waste, in the sense that better, smarter ways of allocating the funding could still see it go to the right people with more research being done overall?

The results from our model suggest the latter. It's not that the money fails to go to the right people: the fourth column of Table 1 shows that even in the most pathological cases there's a positive relationship between talent and success. The problem is simply that too much time gets spent on proposal writing, an activity with direct value in itself that is only worth doing if it leads to funding for future research. Academics, both in our model and most likely in reality, are caught in a kind of tragedy of the commons: their individually rational efforts to write a convincing proposal that gains them a slice of the funding pie lead to an equilibrium in which the research output of the system as a whole goes down.

It might be objected that the MEMORY B heuristic indicates that competitive allocation can be reasonable under the right circumstances. It's true that under some conditions, a competitive allocation system can slightly outperform randomly handing out funding. However, this only happens because less successful agents readily drop out and stay there. In real academic environments the pressure to put in funding applications does not make dropping out easy to do (except for the radical step of changing careers entirely).

Many further variants of the model could be explored (*e.g.*, basing bid evaluation on past performance, or incorporating social learning or coalition formation). Additionally, interviews with academics would help to refine the agent heuristics. However, the basic message is clear: if competition for funding forces agents to allocate ever-increasing levels of their time to bid preparation, then overall system efficiency will suffer.

References

1. Agar, M.: My kingdom for a function: modeling misadventures of the innumerate. Journal of Artificial Societies and Social Simulation 6(3), 8+ (2003), http://jasss.soc.surrey.ac.uk/6/3/8.html
2. Clark, T.: ARC grant process counterproductive (July 2009), *Campus Review*, 14 July
3. Denham, J.: The allocations of the science budget 2008/09 to 2010/11 (December 2007), report of the Department for Innovation, Universities and Skills, UK.
4. Goldsworthy, J.: Research grant mania. Australian Universities Review 50(2), 17–24 (2008)
5. Gordon, R., Poulin, B.J.: Cost of the NSERC science grant peer review system exceeds the cost of giving every qualified researcher a baseline grant. Accountability in Research: Policies and Quality Assurance 16(1), 13–40 (2009)
6. Heath, B., Hill, R., Ciarallo, F.: A survey of agent-based modeling practices (January 1998 to July 2008). Journal of Artificial Societies and Social Simulation 12(4), 9+ (2009), http://jasss.soc.surrey.ac.uk/12/4/9.html
7. Ibaraki, T., Katoh, N.: Resource Allocation Problems: Algorithmic Approaches. MIT Press, Cambridge, MA (1988)
8. Iwasa, Y., Pomiankowski, A., Nee, S.: The evolution of costly mate preferences II. The "handicap" principle. Evolution 45(6), 1431–1442 (1991)
9. Lawrence, P.A.: Real lives and white lies in the funding of scientific research. PLoS Biology 7(9), e1000197+ (2009), http://dx.doi.org/10.1371/journal.pbio.1000197
10. Liefner, I.: Funding, resource allocation, and performance in higher education systems. Higher Education 46(4), 469–489 (2003)
11. Link, A., Swann, C., Bozeman, B.: A time allocation study of university faculty. Economics of Education Review 27(4), 363–374 (2008)
12. Massy, W.F.: Resource Allocation in Higher Education. University of Michigan Press (1996)
13. Mills, D.: An introduction to the UK higher education research funding model. Tech. rep., Economic and Social Science Research Council, UK (2006)
14. Schaffer, A.: America's got science talent: the biomedical research community goes bananas for $200 million in stimulus funding (April 2009), *Slate*, 29 April
15. Zahavi, A., Zahavi, A.: The Handicap Principle: A Missing Piece of Darwin's Puzzle. Oxford University Press (1997)