CENTRALISED FUNDING AND THE DIVISION OF COGNITIVE LABOUR

ABSTRACT. Project selection by funding bodies directly influences the division of cognitive labour in scientific communities. I present a novel adaptation of an existing agent-based model of scientific research, in which a central funding body selects from proposed projects located on an epistemic landscape. I simulate four different selection strategies: selection based on a god's-eye perspective of project significance, selection based on past success, selection based on past funding, and random selection. Results show the size of the landscape matters: on small landscapes historical information leads to slightly better results than random selection, but on large landscapes random selection greatly outperforms historically-informed selection.

Word count: 4359

INTRODUCTION

National funding bodies support much of contemporary science. The selection criteria for funding have gained increasing attention within philosophy of science (Gillies, 2008; O'Malley et al., 2009; Haufe, 2013; Lee, 2015). Meanwhile, there has been growing interest in model-based approaches to understanding the social epistemic activities of scientists (Kitcher, 1990; Strevens, 2003; Weisberg and Muldoon, 2009; Grim, 2009; Zollman, 2010). The current paper builds on previous modelling tools to explore the effects of centralised selection mechanisms on the division of cognitive labour and the ability of scientific communities to efficiently discover significant truths.

Science aims at discovering significant truths, i.e. not just any truths, but truths that will eventually contribute in a meaningful way to well-being (Kitcher, 2001). This is the justification for the public support of science, including basic science (Bush, 1945). Some funding terminology: scientific projects have high *impact* (ex post) if they result in significant truths; projects have high *merit* (ex ante) if they are predicted to have high impact.

Polanyi (1962) analysed merit as being composed of three components: scientific value, plausibility and originality. Polanyi notes an essential tension between plausibility and originality: the more original a project, the more difficult it is to evaluate its plausibility. Polanyi advocates selection by peer review as a conformist position, that sacrifices the occasional meritorious original project while ensuring all supported research projects are plausible, to "prevent the adulteration of science by cranks and dabblers" (p. 8). Gillies (2008, 2014) takes an opposing position, arguing that the cost of losing (infrequent) highly original and meritorious research is much greater than the cost of occasionally supporting implausible research that ends up being of low impact. As an alternative to peer review, Gillies advocates random selection. The tension between plausibility and originality is clearly relevant to questions of effective division of cognitive labour, and has direct links to science policy. This tension, and its complexity, is explored in this paper.

I will argue that the results of the simulations presented are both significant and surprising. The simulations show that, under reasonable parameter values for at least some fields of science, choosing projects at random performs significantly better, in terms of accumulated significant truths, compared to other funding strategies, including project selection by peer review. The results support, to an extent, Gillies' proposal of funding by lottery.

1. Model description

The model explores the influence of different funding mechanisms on the accumulation of significant truths. It builds on the epistemic landscape model developed by Weisberg and Muldoon (2009), extending it by adding representations of centralised funding selection and dynamic changes in project merit. The latter is added to reflect a more realistic picture of scientific merit. For example, Strevens (2003) discusses the effect of a successful discovery on all further pursuits of the same question: they no longer have any merit, as they lose all originality. Several dynamic processes affecting merit are detailed later in the paper.

The model represents a population of scientists exploring a topic of scientific interest. They are all funded by the same central funding body to pursue projects of varying duration, measured in years. Each project's significance is allocated in advance by the modeller, from a "god's-eye" perspective. When grants end scientists successfully complete their project. Their projects' results contribute to the collection of significant truths in the field's corpus of knowledge. Funding mechanisms are compared by their ability to generate this accumulation of significant truths.

For simplicity, scientists in the model (unrealistically) do not share their findings nor explore similar projects during research. They only work on the project for which they were funded and they only share their results at the end of a grant. The social processes set aside here have been explored in previous works (Grim, 2009; Zollman, 2010). Future work may combine the different models towards a unified picture of the division of cognitive labour.

Funding is represented as a process of selection. In every time step, the scientists whose grants have run out are placed in a pool of candidates along with new entrants to the field, and the modelled funding mechanism selects from this pool of candidates those who will receive funding and carry out research projects. Modelled funding mechanisms differ in the way they select individuals, as outlined below.

Actual potential: Actual potential, which can only be known from a god's-eye perspective, is the significance of a project's results *were it successfully completed today*. In the absence of time-dependant merit, actual potential is simply the significance of the project's results. However, in the presence of time-dependence the significance could change between the initiation of the project (at the point of funding) and its completion (at the point of contributing the results to the relevant corpus). This means that in the presence of time-dependence, actual potential might diverge from the eventual contribution of the project.

Estimated potential: Estimated potential is the scientific community's ex ante evaluation (assumed, for simplicity, to be single-valued) of the merit of a proposed project. This prediction is taken to rely on the known contributions of past projects which bear some similarity to the proposed project, and so depends on the history of research projects in the field. In representing decisions based on the research community's prediction, this selection method is akin to peer-review.

Past funding: Under this mechanism, funding is allocated to those scientists who already received funding in the past, and only to them. The model (unrealistically) represents all scientists as being of equal skill, and so this mechanism cannot be taken to mean the selection of the most "intrinsically able" scientists. Rather, this mechanism is included as a "most conservative" option, not admitting any new researchers to the field beyond the field's original investigators.

Lottery: Under a lottery, all candidates have equal chances of being funded. The lottery option serves both as a natural benchmark for other funding methods, and as a representation of the mechanism proposed by Gillies (2014).

The essence of the model is the comparison of the performance of these selection mechanisms in generating results of high significance over time under various conditions.

To represent in the model the time-dependence of merit, the significance contributions of different project results are allowed to change over time as a response to scientists' actions. Three dynamic processes are included in the model (details in §2.5). Two processes involve a reduction of significance following a successful project or breakthrough, which reflects the one-off nature of discovery. The third process involves an increase in significance when a new avenue of research is opened by a significant discovery. Simulations based on the model show that these dynamic processes have a significant effect on the relative performance of different funding strategies.

2. Simulation details

2.1. **Simulating the epistemic landscape.** To investigate the complex nature of the domain being modelled, the model was turned into a computer simulation.[1] The basic structure of the landscape simulation follows Weisberg and Muldoon's, of a two-dimensional configuration space, charted with two coordinates x and y, with an associated scalar field represented in a third dimension as height along the z axis. Each (x, y) coordinate pair specifies a different potential research project; the closer two projects are on the landscape, the more similar they are. The scalar value associated to the coordinate represents the significance of the result obtained on a successful completion of the project, were it completed today (allowing for time dependence). The limit to two spatial dimensions of variation between projects is likely to be unrealistic (Wilkins, 2008), but a higher-dimensional alternative would make the model much less tractable.

In each run of the simulation, the landscape is generated anew in the following process:

(1) Initialise a flat surface of the required dimensions.
(2) Choose a random location on the surface.
(3) Pick random values for relative height, width along x, and width along y.
(4) Add to the landscape a hill at the location chosen in step 2 by using a bivariate Gaussian distribution with the parameters picked in step 3.
(5) Repeat steps 2-4 until the specified number of hills is reached.
(6) Scale up linearly the height of the landscape according to the specified maximum height.

This process generates the "god's-eye" perspective of the research potential of the domain. Here and later, random variables are used to fill-in parameters whose existence is essential for the simulation, but where (1) the specific values they take can vary across a range of valid model targets, and/or (2) there is no compelling empirical evidence to choose a particular value. This requires, however, several runs of the simulation for each configuration, to average out the effects of random variation.

[1]Source code for the simulation is available from the author on request.

2.2. **Simulating agents.** The agents in the model represent scientists investigating the epistemic landscape. Each agent represents an independent researcher or group, and is characterised by its location on the landscape, representing the project they are currently pursuing, and a countdown counter, representing the time remaining until their current project is finished. Like Weisberg and Muldoon's "hill climbers", agents are simulated as local maximisers. Agents follow the following strategy every simulation step:

(1) Reduce countdown by 1.
(2) If countdown is not zero: remain in same location.
(3) If countdown is zero: contribute to the accumulated significance the significance of the current location, and attempt to move to the highest local neighbour.

In the simulation, the agents are identical, in the sense that any agent, when successfully completing a project of a given significance, will contribute exactly that amount to the accumulated significance of the field. This simplification ignores natural ability and gained experience, and stems from a focus on a particular approach to science funding, which funds *projects*, rather than funding *people*. The focus is informed by the explicit policies of certain funding bodies, like the National Institutes of Health (NIH), reflected, for example, in the institution of blind peer review. Thus, the results of the current work would not extend to the minority of science funding bodies, such as the Wellcome Trust, that make explicit their preference to fund people rather than projects.

The *local neighbourhood* of an agent is defined as the 3×3 square centred on their current position. The attempt to move to the highest neighbour depends on the selection (funding) mechanism, as discussed below. The *accumulated significance*, which is the sum of all individual contributions to significance, is stored as a global variable of the simulation and used to compare strategies.

In the beginning of the simulation, a specified number of agents are seeded in random locations on the landscape, with randomly generated countdowns selected from a specified range of values. An example of an initial seeding of agents can be seen in Fig. 1.

In the absence of selection and time-dependence, the course of the simulation is easy to describe: agents begin in random locations on a random landscape, and as the simulation progresses the agents finish projects and climb local hills, until, after an amount of time which depends on the size of the landscape, the number and size of peaks, and the duration of grants, all agents trace a path to their local maxima and stay there. Since agents increase their local significance during the climb, the rate of significance accumulation increases initially, until all agents reach their local maxima, at which point significance continues accumulating at a fixed rate indefinitely. This is the dynamic

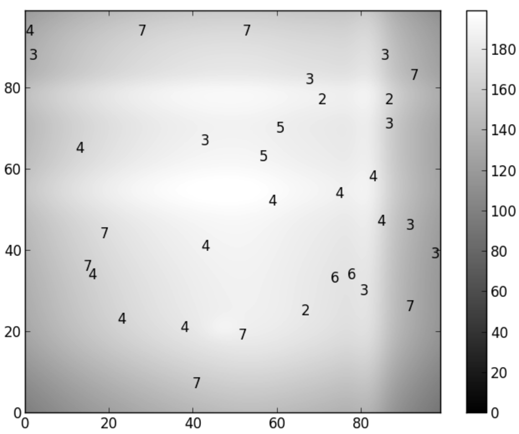

FIGURE 1. Landscape simulation with initial seeding of agents. Each number on the landscape represents an agent at its location, with the value of the number representing the agent's countdown. The colours indicate the height (significance) of each position (project) in the landscape.

seen in Weisberg and Muldoon's simulation for a pure community of "hill climbers", and its unrealistic nature highlights the importance of simulating the time-dependence of significance.

2.3. **Simulating communal knowledge.** In addition to their contribution to significance, agents also contribute to the *visibility* of the landscape (Muldoon and Weisberg, 2011). The visibility of a project represents whether the scientific community, and especially funding bodies, can estimate the significance contribution of that project. Initially, the entire landscape is invisible, representing full uncertainty. Upon initial seeding of agents, each agent contributes vision of their local neighbourhood, as defined above, to the total vision. As the agents move, they add vision of their new local neighbourhood. Visibility is used in the *best_visible* funding mechanism described below.

The simulation represents visibility in a simplistic manner by assigning binary values: either the community knows what the significance of a project will be, or it does not. A more realistic representation will allow partial visibility, with some distance decay effect, such that the community would still be able to make predictions of significance for less familiar projects, but these predictions will have a probability of being wrong, with the probability of error increasing the more unfamiliar these projects are. This addition, however, will be computationally heavy, as it requires maintaining multiple versions of the landscape, both for the real values and for the estimated values.

2.4. **Simulating funding strategies.** The aim of the model is to explore the effects of funding mechanisms on the population and distribution of investigators. Since the aim is to simulate current funding practices (albeit in a highly idealised manner), and since current funding practices operate in passive mode (choosing from proposals originating from scientists rather than dictating which projects aught be pursued), the guiding principle of the simulation is that a funding mechanism is akin to a selection process: at each step of the simulation, the actual population of agents is a subset of the candidate or potential population, where inclusion in the actual population follows a certain selection mechanism.

Funding mechanisms are simulated in the following manner:
Every step:

(1) Place all agents with zero countdown in a pool of "old candidates".
(2) Generate a set of new candidate agents, in a process identical to the seeding of agents in the beginning of the simulation.
(3) Select from the joint pool of (old candidates + new candidates) a subset according to the selection mechanism specified by the funding method.
(4) Only selected agents are placed on the landscape and take part in the remainder of the simulation, the rest are ignored.

The simulation can represent four different funding mechanisms:

best: selects the candidates which are located at the highest points, regardless of the visibility of their locations. This simulates a mechanism which selects the most promising projects from a god's eye perspective. This overly optimistic mechanism does not represent a real funding strategy. Rather, it serves as an ideal benchmark against which realistic funding mechanisms are measured.

best_visible: filters out candidates which are located at invisible locations, i.e. candidates who propose to work on projects which are too different from present or past projects. It then selects the candidates in the highest locations from the remainder. This strategy is closer to a realistic representation of selection by peer review. Note that even this version is epistemically optimistic, as it assumes the selection panel has successfully gathered all available information from all the different agents, both past and present.

lotto: selects candidates at random from the candidate pool, disregarding the visibility and height of their locations.

oldboys: represents no selection: old candidates continue, no new candidates are generated.

The key parameters for all funding mechanisms are the size of the candidate pool and the size of the selection pool. The size of the candidate pool, which in turn depends on the size of the new candidate pool (as the size of the old candidate pool emerges from the simulation), has been chosen in the simulations such that the total candidate pool is equal in size to the initial number of agents (except *oldboys* where there are no new candidates). This means the success probability changes between funding rounds, around a mean which is equal to 1/(average countdown). With an average grant duration of five years, this yields a success rate of 20%, close to the real value in many contemporary funding schemes (NIH, 2014). The number of grants awarded each year is set to equal the number of grants completed each year, maintaining a fixed size for the population of investigators.

For simplicity, the simulated funding mechanisms do not take into account the positions of existing agents on the landscape, except indirectly when considering their vision. Future simulations may consider a selection mechanism which explicitly favours either diversity or agglomeration, though one expects difficulties in operationalisation and measurement of epistemic diversity.

2.5. **Simulating merit dynamics.** To make the simulation more realistic, the significance of projects is allowed to change over time in response to research activities of the community of investigators. Three such dynamic processes are included in the simulation:

Winner takes it all: As was made explicit by Strevens (2003), the utility gain of discovery is a one-off event: the first (recognised) discovery of X may greatly contribute to the collective utility, but there is little or no contribution from further discoveries of X. In the simulation, this is represented by setting the significance of a location to zero whenever an agent at that location has finished their project and made their contribution to accumulated significance. This effect is triggered whenever any countdown reaches zero, which makes it quite common, but it has a very localised effect, only affecting the significance of a single project.

Reduced novelty: When a researcher makes a significant discovery, simulated by finishing a project with associated significance above a certain threshold, the novelty of nearby projects is reduced, which in the model is simulated by a reduction of significance in a local area around the discovery.

New avenues: When a researcher makes a significant discovery, it opens up the possibility of new avenues of research, simulated in the model by the appearance of a new randomly-shaped hill at a random location on the landscape.

3. Results and discussion

Here I present the results of simulations of different setups of interest, exploring the relative success of different funding mechanisms under different conditions.

All simulation results show a comparison between the four funding mechanisms, as a plot of total accumulated significance (arbitrary units) at the end of the simulation run, averaged over five runs with different random seeds. In all simulations the range of countdowns was 2 to 7. The number of individuals was set to equal (size of landscape)$^{3/4}$. Simulations were ran for 50 steps. The trigger for significance-dependant processes was 0.7 of the global maximum. Results are shown for a small landscape (50×50) in Fig. 2 and for a large landscape (500×500) in Fig. 3.

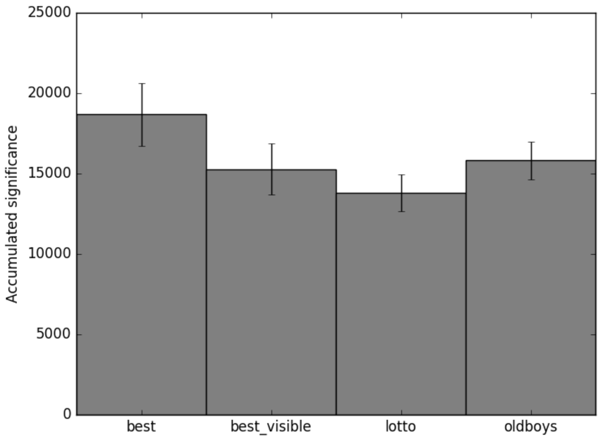

Figure 2. Comparison of significance accumulation under different funding mechanisms, small landscape (50×50).

To get a feeling for how the community is affected by the funding mechanism, I present visualisations of the state of the landscape at the end of the simulation run for the two funding mechanisms mentioned in the introduction (*best_visible* and *lotto*) in Fig. 4. Note that due to the *winner takes it all* dynamic process it is possible to "see" the past trajectory of exploration, as completed projects leave behind highly localised points of zero (remaining) significance. This allows for a visual representation of the division of cognitive labour that emerges under different funding schemes.

As is clear from the simulations, the *best* funding mechanism is indeed best at accumulating significance over time, though with various lead margins over the second best strategy. In the presence of dynamic

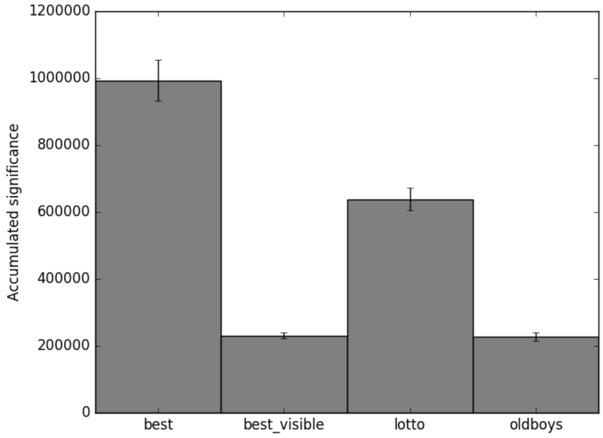

FIGURE 3. Comparison of significance accumulation under different funding mechanisms, large landscape (500×500).

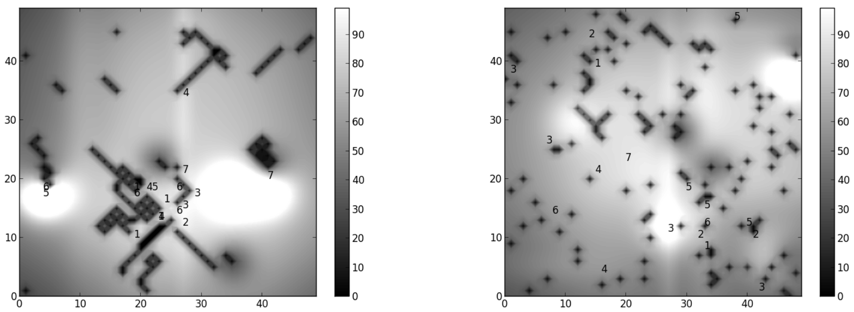

(A) *best_visible* (B) *lotto*

FIGURE 4. Landscape visualisation at the end of the simulation run under different funding mechanisms.

processes, *best* is in the best position to locate new avenues for research, wherever they show up. However, as mentioned above, the *best* funding strategy is not realisable, as it requires a god's eye view of the epistemic landscape.

On the small landscape the three strategies, *best_visible*, *oldboys*, and *lotto* perform roughly similarly, with *lotto* at a small disadvantage as it cannot make use of valuable information from past successes. It seems counter-intuitive that *best_visible* performs worse than *oldboys*. A possible explanation is the effect of reduced novelty: *best_visible* tends to cluster scientists around the most promising projects, and so when one makes a breakthrough it reduces the significance of contributions for all groups working on similar projects (the phenomenon known in contemporary science as "scooping"). This excessive clustering around fashions is not present in *oldboys* or *lotto*.

On the large landscape *lotto* greatly outperforms *best_visible* and *oldboys*. This is because new avenues on a large landscape are likely to spawn outside the visibility of the agents, where *lotto* can access them but the other two strategies cannot. In the smaller landscape this effect is not apparent, as the relative visibility is larger, and therefore the chance of a new avenue appearing within the visible area is larger.

Conclusion

This paper presented a way to extend existing epistemic landscape models so that they can represent selection by a central funding body and time dependence of significance. This model was used in computer simulations to compare the effectiveness of different idealised versions of selection criteria, most notably selection based on past successes (akin to peer review), random selection and no selection. The most significant result from the simulation was that on a large landscape, when a topic can be explored in many ways that could be very different from each other, random selection performs much better than selection based on past performance.

This result fits in with a general result from the body of works on agent-based models of scientific communities, that shows diversity in the community trumps individual pursuit of excellence as a way of making communal epistemic progress. The tension of science funding, between originality and plausibility, is thus a part of the broader tension between diversity and excellence, between exploration and exploitation.

Previous social epistemology models have focused on the role of *internal* factors in shifting the balance between exploration and exploitation. Kitcher (1990); Strevens (2003) look at reward structures (of internal credit, not external monetary rewards) and individual motivation towards credit or truth. Grim (2009); Zollman (2010) look at information availability and information transfer between scientists, and at individual beliefs. Weisberg and Muldoon (2009) look at individual researchers' social strategy: follower or maverick.

The current work is the first within this modelling lineage to look at the effects of an *external*, *institutional* factor: selection by a centralised funding body. The current paper brings this line of research closer to having a direct relevance to science policy. Hopefully future work in this vain will continue this trend, to deliver on the challenge set out by Kitcher (1990, p. 22):

> How do we best design social institutions for the advancement of learning? The philosophers have ignored the social structure of science. The point, however, is to change it.

We could start by advocating for funding mechanisms that allow for more exploration.

References

Bush, V. (1945). *Science, the endless frontier: A report to the President*. Washington: U.S. Government printing office.

Gillies, D. (2008). *How should research be organised?* London: College Publications.

Gillies, D. (2014). Selecting applications for funding: why random choice is better than peer review. *RT. A Journal on research policy and evaluation 2*(1).

Grim, P. (2009). Threshold phenomena in epistemic networks. In Complex adaptive systems and the threshold effect: Views from the natural and social sciences: Papers from the AAAI Fall Symposium, pp. 53–60.

Haufe, C. (2013). Why do funding agencies favor hypothesis testing? *Studies in History and Philosophy of Science Part A 44*(3), 363–374.

Kitcher, P. (1990). The division of cognitive labor. *The journal of philosophy 87*(1), pp. 5–22.

Kitcher, P. (2001). *Science, truth, and democracy*. New York: Oxford University Press.

Lee, C. J. (2015). Commensuration bias in peer review. *Philosophy of Science 82*(5), 1272–1283.

Muldoon, R. and M. Weisberg (2011). Robustness and idealization in models of cognitive labor. *Synthese 183*(2), 161–174.

NIH (2014). Success rates - NIH research portfolio online reporting tools (RePORT). `http://report.nih.gov/success_rates/`, Accessed 11 July 2014.

O'Malley, M. A., K. C. Elliott, C. Haufe, and R. M. Burian (2009). Philosophies of funding. *Cell 138*(4), 611–615.

Polanyi, M. (1962). The republic of science: Its political and economic theory. *Minerva 1*, 54–73.

Strevens, M. (2003). The role of the priority rule in science. *The journal of philosophy 100*(2), 55–79.

Weisberg, M. and R. Muldoon (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of science 76*(2), 225–252.

Wilkins, J. S. (2008). The adaptive landscape of science. *Biology and philosophy 23*(5), 659–671.

Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis 72*(1), 17–35.