Centralised Funding and Epistemic Exploration

Shahar Avin

Penultimate draft. Final draft forthcoming in *The British Journal for the Philosophy of Science (BJPS)*.

Abstract

Computer simulation of an epistemic landscape model, modified to include explicit representation of a centralised funding body, show the method of funding allocation has significant effects on communal trade-off between exploration and exploitation, with consequences for the community's ability to generate significant truths. The results show this effect is contextual, and depends on the size of the landscape being explored, with funding that includes explicit random allocation performing significantly better than peer-review on large landscapes. The paper proposes a way of incorporating external institutional factors in formal social epistemology, and offers a way of bringing such investigations to bear on current research policy questions.

1 Introduction

An important topic in social epistemology is the efficient trade-off between exploration (of new theories, methods, and so forth) and exploitation (of same), seeking both to resolve what might be an optimal division of cognitive labour, and how such a division might be achieved (Goldman and Blanchard [2016]). With too little exploration scientists are restricted to selecting the best theory of a bad lot, never hitting on that most simple and fruitful world view; with too much exploration, science is swamped by poorly-vetted contenders for our attention, wasting limited resources that could be put to better use in advancing known theories and methods.

Previous social epistemology models have focused on the role of internal factors in shifting the balance between exploration and exploitation. Kitcher ([1990]) and Strevens ([2003]) look at reward structures (of internal credit, not external monetary rewards) and individual motivation towards credit or truth. Grim ([2009]) and Zollman ([2010]) look at information availability and information transfer between scientists, and at individual beliefs. Weisberg and Muldoon ([2009]), Alexander *et al.* ([2015]), Thoma ([2015]), and Pöyhönen ([2016]) look at individual researchers' social learning strategies.

Current models are limited because they do not account for the influence of external factors on the exploration/exploitation trade-off. One important external source of direct influence is the funding allocation mechanism: for academic research supported by grant peer review, central funding bodies play a key role in determining which research projects are carried out, and thereby have a direct influence on the division of cognitive labour. The effects of different funding mechanisms on the exploration/exploitation trade-off in science is already discussed amongst scientists, in the context of research policy (Brezis [2007]; Graves *et al.* [2011]; Herbert *et al.* [2013]; Bollen *et al.* [2014]; Fang and Casadevall [2016]).

This paper brings the modelling techniques of formal social epistemology to bear on the role of centralised funding bodies, connecting the modelling literature to the literature on science funding, including from within philosophy of science (Polanyi [1962]; Gillies [2008], [2014]). The current work is the first within the modelling lineage outlined above to look at the effects of an external, institutional factor. By modifying the model introduced by Weisberg and Muldoon ([2009]) to allow representation of science funding policies, the current paper brings formal social epistemology closer to having a direct relevance for research policy.

I will argue that the results of the simulations presented based on the modified model are both significant and surprising. Not only do funding mechanisms matter significantly for the trade-off between exploration and exploitation, but the simulations show that, under reasonable parameter values for at least some fields of science, introducing a random element into the process of project selection leads to better performance, in terms of accumulated significant truths, compared to funding strategies that rely on historical information, including project selection by peer review. This is because the role of so-called 'mavericks' or 'explorers' becomes more important when there is more landscape to explore relative to the size of the community, and grant peer review biases funding towards more conservative so-called 'followers' or 'extractors'. On the social epistemology side, the results show a different way to achieve an effective trade-off between exploration and exploitation, one that is more nuanced to the context of a specific field of research. On the science funding side, the results support, to an extent, the proposal by Gillies and others of explicitly introducing randomness into the selection process of research projects.

2 Theoretical background

Social epistemology of science has been wrangling for some time with the question of the efficient trade-off between exploration and exploitation. Flexing our whiggish muscles we could map the distinction between Normal and Revolutionary science onto this problem (Kuhn [1970]), but the contemporary, model-driven, economic-language flavoured form of the problem comes from (Kitcher [1990], [1993]). Kitcher's models use the transition from individual pursuit of truth to collective pursuit of truth to show how non-epistemic motives (for example the pursuit of credit) can lead to better collective epistemic results – specifically by shifting the above-mentioned trade-off towards more exploration. One line of further work stayed close to Kitcher's marginal utility models, but gradually refined the factors that go into the credit maximisation calculations performed by the scientists in the model, taking into account subtleties of competition and collaboration, and their effect over time on the population-level trade-off between exploration and exploitation (Brock and Durlauf [1999]; Strevens [2003]; De Langhe [2014]). A second line looked beyond credit maximisation to the interaction between scientists as a factor that affects the division of labour: one sub-branch explored models with a homogenous population of scientists but with varying levels of connectivity within the population (Grim [2009]; Zollman [2010]), while another sub-branch explored models with heterogenous populations, with some scientists being more exploration-oriented than others (Weisberg and Muldoon [2009]; Alexander *et al.* [2015]; Thoma [2015]; Pöyhönen [2016]). An important methodological innovation of the latter branch is the introduction of epistemic landscapes, which explicitly add to the model something for the scientists to explore.

The works surveyed above have given us sophisticated and helpful advice on how to get a better balance of exploration and exploitation in science: one should

- divide credit equally between contributors,

- provide credit only to the first discoverers of a breakthrough,
- limit connectivity between scientific groups, and
- provide extra rewards for mavericks.

There is, however, a pernicious shared assumption across all the works above: it is the scientists themselves, through their individual choices and guided by constraints and incentive structures, who get to decide what research gets done. This is at best only partially true. There are more potential research projects and more aspiring researchers than can be supported (financially, institutionally) at any given time, and so the selection of the (rather small) subset of research projects that get carried out at any given time is largely under the control of centralised funding agencies (at least if we limit the scope to academic science). Now, this is not to say scientists have no control over their research projects; indeed, this degree of freedom is often the reason given for why talented and educated individuals are willing to suffer much lower salaries than their colleagues employed in industry or running private ventures. The freedom of the researcher is, however, limited: she gets to shape her research proposal, but must strike a balance between what she thinks would be interesting or good for her career and what might get funded, and even then more often than not her proposal will not be funded and the project will not take place, thus mattering little to the collective division of labour. It should be mentioned, though, that once funded the researcher has much more control over her contribution to the collective pursuit of knowledge, at least until the grant runs out; in this time-limited domain the models surveyed above would take full force, but they should be adjusted accordingly. More importantly for our investigation here, the role of centralised funding bodies and the effects of their selection processes on the division of cognitive labour should at least be considered if we want a more complete picture of the exploration/exploitation trade-off.

National funding bodies support much of contemporary science (Kennedy [2012]). The mechanisms of selection for funding, primarily peer review of prospective research-project proposals (“grant peer review”), and the selection criteria deployed within these mechanisms, have gained increasing attention within philosophy of science outside formal social epistemology (Gillies [2008]; O'Malley *et al.* [2009]; Haufe [2013]; Lee [2015]). To start bridging the conceptual gap between formal social epistemology and studies of science funding we should introduce some of the concepts used in the latter literature. A good common starting ground is the view that science aims at discovering significant truths, namely not just any truths, but truths that will eventually contribute in a meaningful way to wellbeing (Kitcher [2001]). This is the justification for the public support of science, including basic science (Bush [1945]). Some funding terminology: research 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, as do Greenberg ([1998]) and Fang and Casadevall ([2016]). We thus find the exploration/exploitation tension at the heart of philosophical analysis of science funding mechanisms, and are ready to deploy the tools perfected within formal social epistemology to explore more carefully the relevance of funding mechanisms to the question of the division of cognitive labour. To do so, I will present yet-another modified version of Weisberg and Muldoon's epistemic landscape model, but one that has explicit representation of centralised funding.

3 Model description

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 (this is the 'height' of each point of the landscape).[1] When grants end scientists successfully complete their projects. 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.

We should note two ways in which this already deviates from the models in existing literature. One clear deviation is the addition of a funding mechanism, which means scientists do not simply persist in the model throughout the simulation: sometimes they get renewal grants and can continue previous research programmes, but sometimes their funding runs out and they are eliminated from the simulation; meanwhile, new entrants to the field appear and compete for funding. A more subtle deviation is the time duration represented in the model. One 'turn' in previous models represented a timescale of a few months, whereas one turn in my simulation is one year. While this is on roughly the same order of magnitude, it would have an effect on reasonable assumption about how much work a scientist can do per turn, and also how coarse or fine-grained we should take our model coordinates to be, which influences reasonable movement strategies (how 'far' would a scientist go when transitioning from one project to the next).

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. This is a clear deviation from previous work that has focused on agent-agent interaction as a dominant factor in division of cognitive labour. By setting aside the direct interaction factor I show that even when only considering indirect interaction (via the centralised funding body and via the landscape itself, as described below) there is a significant effect on the division of cognitive labour. I look

[1]The assigned value of significance is not meant to represent an objective value of the research project, nor argue that such assignment indicates such objective value exists. Rather, the assigned significance represents the intersubjective evaluation of the research project's epistemic contributions were it completed, given the notions of wellbeing prevalent in the society in which the research would take place.

forward to future work that will consider the interplay of these factors.

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, 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 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 is a method of allocating funding 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. This funding mechanism reproduces the model dynamics of previous work: the same population persists throughout the simulation.

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 an idealised representation of the mechanisms proposed by Greenberg, Gillies, Fang and Casadevall.

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 §4.5). Two processes involve a reduction of significance following a successful project or breakthrough, which reflects the one-off nature of discovery (a mechanism similar to one of these has been used by Pöyhönen in his model). The third process involves an increase in significance when a new avenue of research is opened by a major discovery. Simulations based on the model show that these dynamic processes have a significant effect on the relative performance of different funding strategies.

4 Simulation details

4.1 Simulating the epistemic landscape

Like previous work, the model was explored using computer simulations.[2] The basic structure of the landscape simulation follows prior work by representing a research topic with a range of possible research projects as 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]),

[2]Source code for the simulation is available at <https://github.com/shaharavin/FundingSimulator>. Since my model introduces major changes, new code was written instead of reusing previously existing code. Nonetheless, I would encourage future work to rely on existing code, especially when available as open-source.

but a higher-dimensional alternative would make the model much less tractable.[3]

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. The control over the number and size of hills allows variation of 'ruggedness', and values used in simulation runs have tended towards higher ruggedness, as least when compared to Weisberg and Muldoon's model and similar later iterations; this is in accordance with the model presented by Pöyhönen.

[3]A known limitation of low-dimensionality is an artificial amplification of 'peak shift' problems, as transition between distant peaks is easier in higher dimensions (Gavrilets [2004]). This is a problem for any landscape model that features hill-climbers, including this work and previous work done on epistemic landscapes. The negative effect is much reduced, however, if agents are allowed to 'hop', as in the models of Grim ([2009]) and Thoma ([2015]), or if the population itself can change, for example in response to funding decisions as presented here.

4.2 Simulating agents

The agents in the model represent scientists investigating the research topic. Each agent represents an independent researcher or small 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.

Note that the update rule, or search strategy, is naive compared to the more advanced strategies explored by Weisberg and Muldoon, Alexander et al., Thoma and Pöyhönen. This is because the focus of the current model is on the effects of centralised funding rather than the effects of researcher strategies. This naive strategy is, however, somewhat realistic if we interpret the coordinate space as being fairly coarse-grained, and so moving far from one's current position would represent a significant (and costly) change to one's lab equipment and practices. Furthermore, due to the fact that agents spend multiple turns in the same location, and the occasional elimination of agents from the landscape, agents tend to have much shorter 'trajectories' on the landscape, and so the specific project-choice strategy they pursue matters less.

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 in the field. This simplification, which also exists in all previous models mentioned in the introduction, ignores natural ability and gained experience. In the context of the current model this can be reasoned as representing a particular approach to science funding, which funds projects, rather than funding people. Focusing on this approach to funding is informed by the explicit policies of certain funding bodies, for example 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 (a similar measure has been proposed by Pöyhönen as superior to the success measures used by Weisberg and Muldoon).

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.[4] 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 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.

[4]The seeding of agents at random locations is likely an unrealistic simplification, as the history of the field, individual training, and social networks are likely to affect the kinds of projects researchers would consider pursuing. These internal factors, all likely to be relevant to diversity and epistemic success, are set aside in the current work as it focuses on the role of the external funding mechanism. Future work will hopefully combine these two streams.

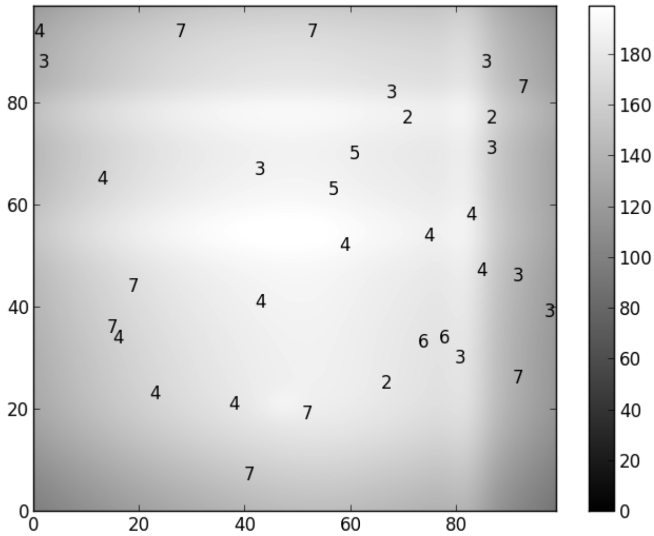

Figure 1: Landscape simulation with initial seeding of agents. Each number on the landscape represents an agent at its location, with the value of the number representing the agent's countdown. Brightness indicates the height (significance) of each position (project) in the landscape. Each simulation run a new landscape is generated using the method described in the text.

4.3 Simulating communal knowledge

In addition to their contribution to significance, agents also contribute to the 'visibility' of the landscape.[5] 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 region to the total vision. As the agents move, they add vision of their new local region. Visibility is used in the 'best_visible' funding mechanism described below.

As one anonymous reviewer pointed out, the vision distance of agents may end up playing a crucial role in the relative success of funding mechanisms, and so the choice of parameters should be explored carefully. Distance between adjacent coordinates, as in previous epistemic landscape models, is a measure of similarity between projects. To get a better handle on how far researchers should be able to 'see', we need a better understanding of just how similar two

[5]Visibility as a model feature has been explored in a different model by Muldoon and Weisberg ([2011]).

adjacent coordinates are. One way to do so is to relate the 'vision' distance to the 'movement' distance. Agent movement was described above as local hill climbing, with a movement distance of one tile. The justification for this movement rule, which also appears in previous modelling work, was that agents are constrained by existing laboratory equipment and practices from changing their research agenda too drastically. However, even given these inertial constraints there should be room to manoeuvre, and so it follows that landscape coordinates are coarse-grained with respect to project similarity, such that a one-tile distance represents a sizeable change in project content.

The movement of agents also allows us to link the theoretical landscape coordinates to empirical results. Boudreau *et al.* ([2016]) measured the correlation between the novelty of research proposals and their peer-review scores. They measured project novelty by the overlap between key terms describing the proposal (assigned by an independent expert classifier) and the terms appearing in existing literature. They found a clear pattern, where proposals score highest when about half their key terms overlap with past literature and half are novel; a precipitous drop in scores occurs when proposals have more than 70% novel terms (Boudreau *et al.* [2016], Fig. 3). If we assume scientists are aware of this pattern, we would expect progressive research proposals from the same researcher to have about 50% novel terms, or perhaps a little less, say 40% or 30% novelty, if inertia is taken into account. The existing movement rule then suggests a one-tile distance on the landscape represents a 50%-70% content similarity. This would make a distance of two tiles represent 25%-50% similarity, and a distance of three tiles represent 6%-25% similarity. Given Boudreau et al's findings, this would make proposals at distance of three or more tiles from past projects effectively un-fundable under peer-review. A high degree of myopia (short vision distance) seems realistic. To account for the difference between movement distance and vision distance, the vision update rule has been set such that agents can 'see' two tiles away: twice as far as they can move.

Boudreau et al's findings provide an interpretation of the model's 'vision distance' in terms of anti-novelty bias. It follows that the representation of visibility in the model as a binary value (either the community can 'see' a potential project or it cannot) is simplistic. 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 contain a mixture of an uncertainty element (a random error due to the novelty of the project) and an anti-novelty bias element. This addition, however, will be computationally costly, as it will require maintaining multiple versions of the landscape, both for the real values and for the estimated values, and so it is left for future work.

4.4 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 ought 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 simulation 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, namely 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 (and in some cases quite a bit higher than) 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, which would have the most direct effect on the question of division of cognitive labour explored here, though it is not clear whether an actual funding scheme could be put together that would successfully operationalise and measure diversity (Stirling [2007]).

4.5 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. A similar dynamic has been introduced by Pöyhönen, though in his model only a fraction of significance is removed. The variation in timescales accounts for this difference: while it is unreasonable that a scientist would 'extract' all the significant results in a single few-month pass over a field, she is much more likely to do so after exploring it for several years.

Reduced novelty: When a researcher makes a significant discovery, simulated by finishing a project with associated significance above a certain threshold, the novelty of similar projects is reduced. Following Polanyi, who considered novelty to be one of the components of scientific merit, the effect of reduced novelty is represented 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.[6]

---

[6]Without going into the details for lack of space, an example of this kind of effect can be seen in the influence of Chargaff's experiment on the perceived value of the experiment of Hershey and Chase, which in turn had an effect on the significance of the work of Watson and Crick (Allen [1975]).

5 Results and discussion

Here I present the results of simulations in a series of experiments, exploring the relative success of different funding mechanisms under different conditions.

In the first four experiments, simulation results show a comparison between different funding mechanisms as a plot of total accumulated significance (arbitrary units) at the end of the simulation run, averaged over five runs per funding mechanism with different random seeds. In all simulations the range of countdowns was 2 to 7, and simulations were ran for 50 steps, simulating 50 years of research with grant durations ranging from two to seven years.

5.1 Experiment 1: Only 'Winner takes it all'

The first experiment starts close to prior work: only one dynamic process is simulated: 'Winner takes it all', the elimination of remaining significance in a very localised region around completed projects. This dynamic process is similar to the one introduced in Pöyhönen's model. The main difference from prior work is the introduction of funding mechanisms, instead of social learning, as the difference maker between experiment groups. Since preliminary results showed an important role for the size of the landscape, results are shown for a range of landscape sizes (Fig. 2). For reference, previous work following Weisberg and Muldoon's model uses the original size of (100×100). The number of peaks and the number of agents were scaled with the size of the landscape, so that the per-tile frequency of peaks and agents was kept fixed (at $1/100$ and $1/200$ respectively).

5.2 Experiment 2: All dynamic mechanisms

The second experiment is identical to the first except each simulation now also incorporates the 'Reduced novelty' and 'New avenues' dynamic mechanisms, which are triggered only when a project of sufficiently high significance is completed. The trigger for significance-dependant processes was set to 0.7 of the global maximum. Results are shown in Fig. 3.

To get a sense 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

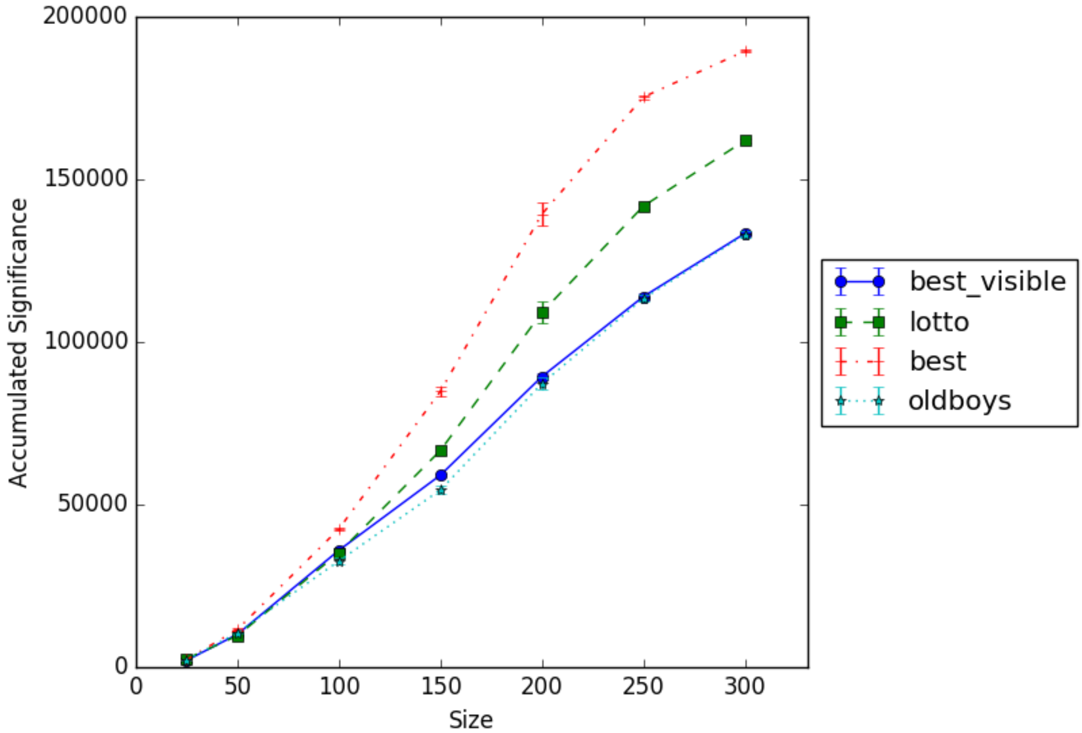

Figure 2: Comparison of significance accumulation under different funding mechanisms for different landscape sizes. Only 'Winner takes it all' dynamic mechanism is at play in the simulations.

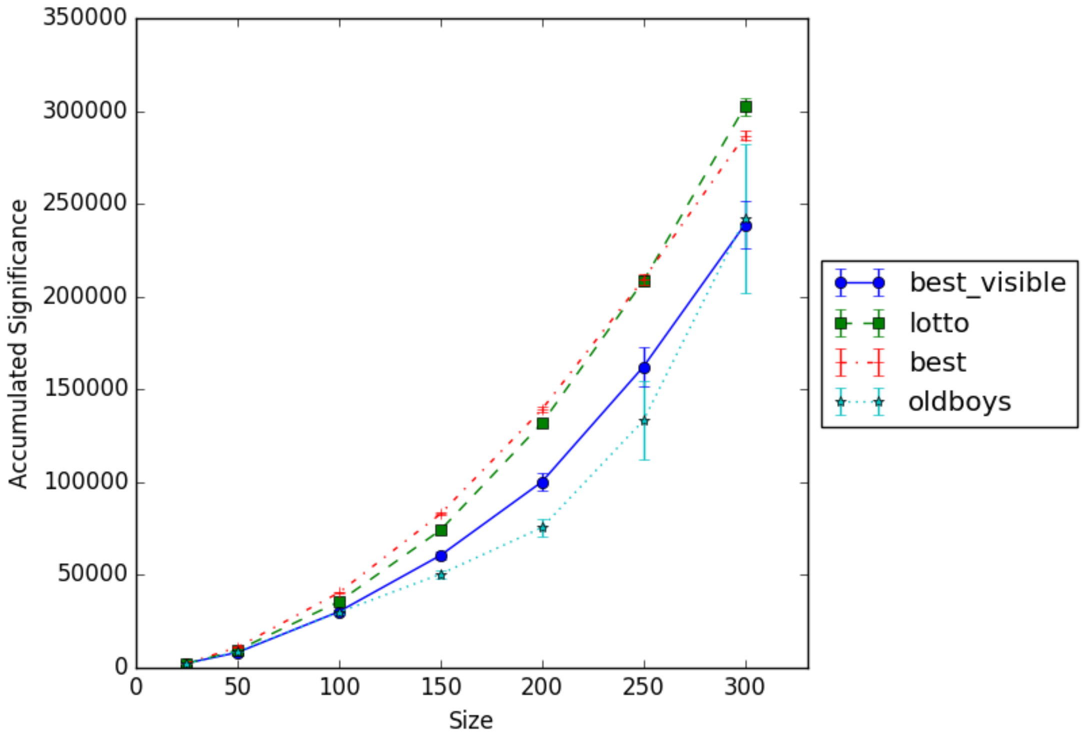

Figure 3: Comparison of significance accumulation under different funding mechanisms for different landscape sizes. All three dynamic mechanisms are at play in the simulations.

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.

5.3 Experiment 3: Adding a new funding mechanism – ‘triage’

One possible explanation for the relative success of the ‘lotto’ funding mechanism in the previous experiments is that it allows the community to access previously ‘unseen’ portions of the landscape, in effect encouraging more exploration than ‘oldboys’ or ‘best_visible’. However, ‘lotto’ ignores knowledge of expected significance gains even when these are known, potentially resulting in waste. This suggests an improvement in the form of a novel funding mechanism, ‘triage’, that provides some funding according to available data within the vision range of the centralised funding body, and distributes the rest of the funding to

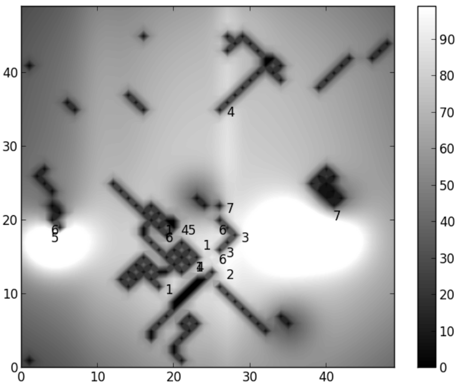

(a) best_visible

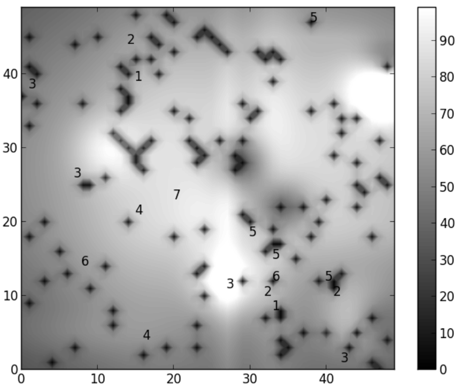

(b) lotto

Figure 4: Landscape visualisation at the end of the simulation run under different funding mechanisms. Brighter locations represent points of higher (remaining) significance).

researchers outside its vision range (these may be considered 'mavericks' in relation to previous work).[7] Something like 'triage' seems to be at play in some contemporary funding agencies who allocate most funds according to merit-based grant peer review, but allocate some funds for high-risk research (though this too is often done using grant peer review, albeit with a different guidance to reviewers). In the simulation I model this mechanism with a 50%/50% divide between 'best_visible' within-vision and 'lotto' outside-vision. Results for simulations with only 'Winner takes it all' are shown in Fig. 5a and for simulations with all three dynamic process in Fig. 5b.

5.4 Experiment 4: Varying the degree of myopia

As discussed above, the vision distance of the funding body is likely to play an important role in the relative success of different funding mechanisms. We have already seen above, from performance of the 'best' funding mechanism, that under full vision selection by merit performs best. §4.3 argued, from comparison to movement distance, that a vision range of two tiles is realistic. It is still interesting to explore just how much vision is required for merit-based selection to outperform funding by lottery. To test this, I ran the simulation with increasing values of 'vision range', under different scenarios.[8] Fig. 6 shows a comparison of accumulated significance at the end of the run across different vision ranges and different funding mechanisms, on a (150×150) landscape with all three dynamic processes.

A vision range of two to three tiles, which seems plausible given the coarse-grained interpretation of the coordinates, leads to an advantage for 'lotto' over 'best_visible'. Further relaxation of the funders' myopia may challenge the significance of some of the results obtained in previous experiments, though notice that 'triage', which incorporates an explicit element of random selection, maintains an edge over 'best_visible' even if vision range is increased to four or five tiles.

---

[7]There are similarities between the 'triage' mechanism and the 'focal randomisation' mechanism proposed by Brezis ([2007]).

[8]In the simulation, whenever a new project is started, the area that is added to the region marked as 'visible' is a circle of radius $r + 1/2$ centred on the new project, where r is the vision range.

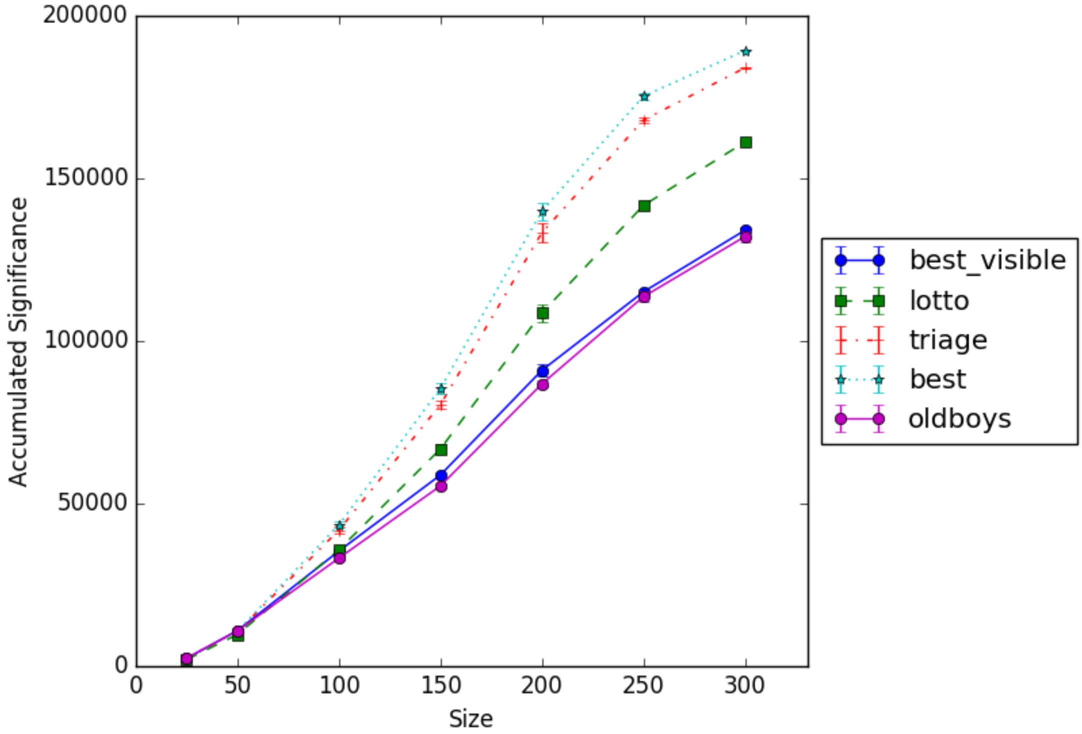

(a) Only 'Winner takes it all'.

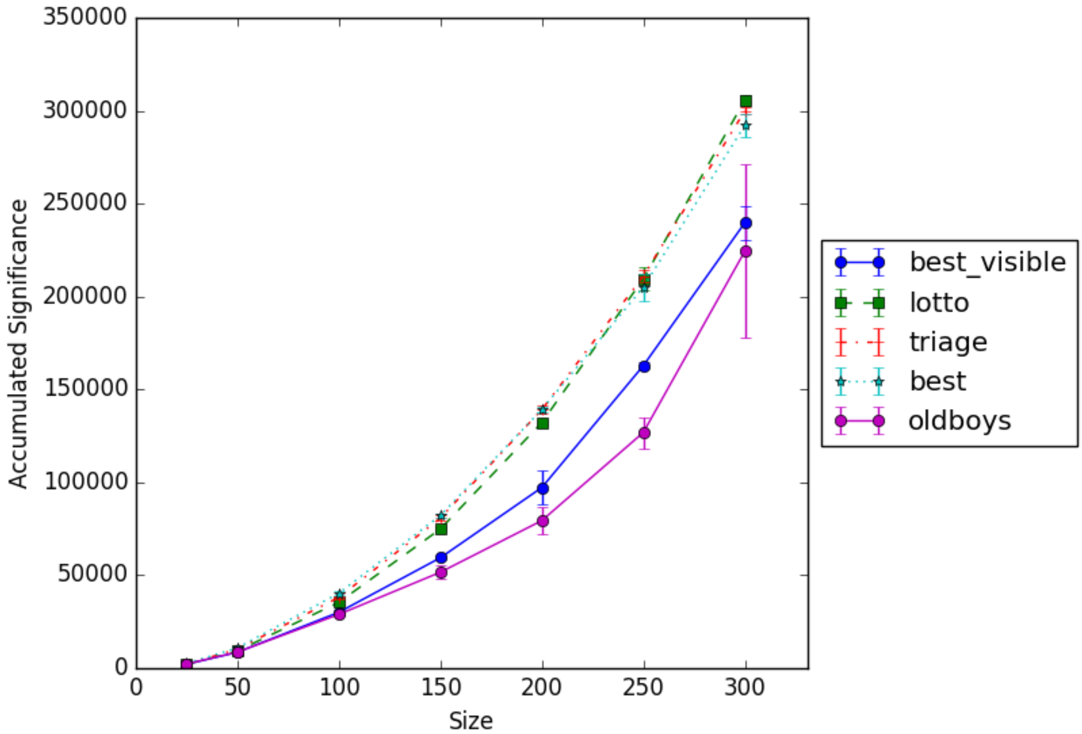

(b) All three dynamic mechanisms.

Figure 5: Comparison of significance accumulation under different funding mechanisms, including 'triage', for different landscape sizes.

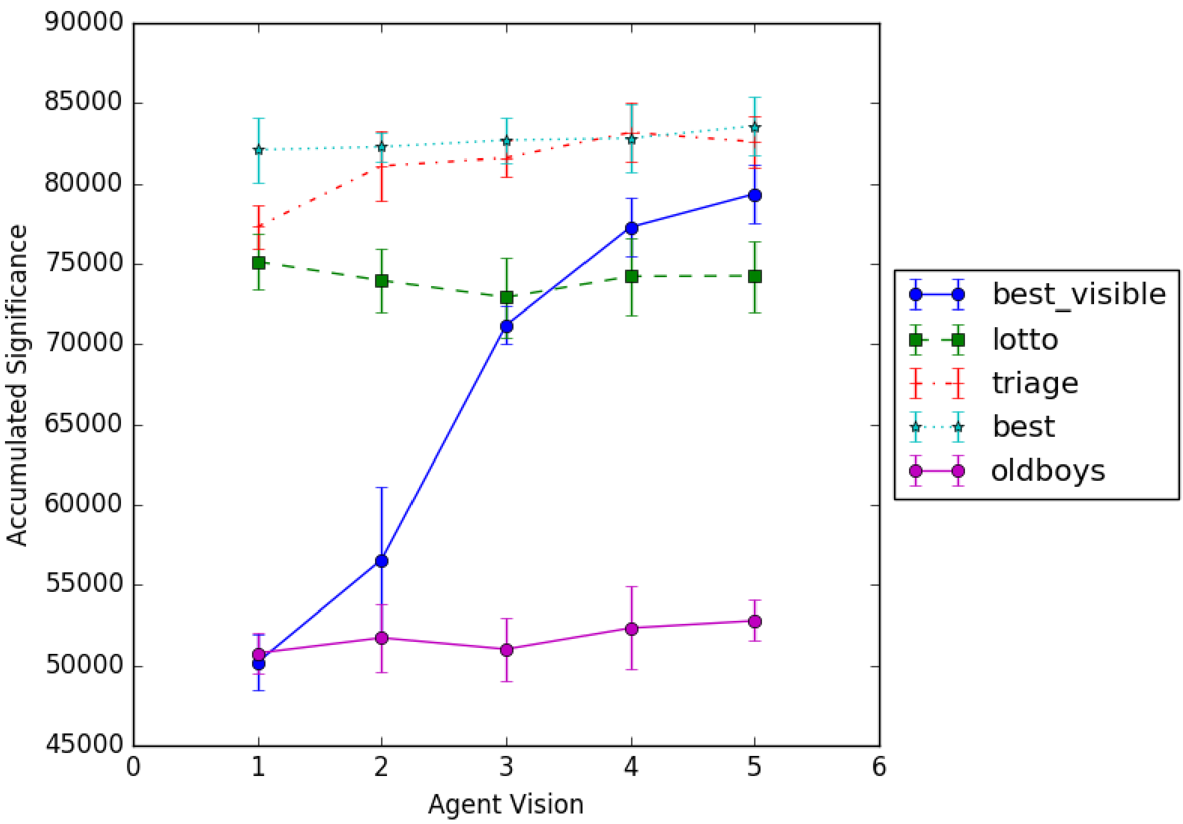

Figure 6: Comparison of significance accumulation under different funding mechanisms for different vision ranges.

5.5 Experiment 5: Variability of individual epistemic gain

So far, experiments have looked at the accumulated significance at the end of the simulation, after 50 simulation steps, representing 50 years of research in the domain. This was considered an appropriate measure from the perspective of evaluating the social epistemic effectiveness of different funding mechanisms under different conditions. However, we may worry that funding bodies are not interested merely in long term epistemic gains, but also in year-by-year, individual-by-individual epistemic gains, which are required to justify to other branches of government the need for continued support for research.[9] We can imagine, for example, that while funding by 'lottery' or 'triage' may outperform 'best_visible' in the long run, they may have higher individual variability in terms of epistemic contributions which would be considered too risky for funders. To test this, I've compared the discrete epistemic contributions of individuals across the 50 years represented in the simulation, across the different funding mechanisms, using a box-and-whisker plot to demonstrate the distribution of values. Results are presented in Fig. 7 for a (150 × 150) landscape. In the presence of only the 'winner takes it all' dynamic mechanism (Fig. 7a) 'best_visible' does indeed provide more of a minimum-quality guarantee than 'lotto', but in the presence of all three dynamic mechanisms (Fig. 7b) this safeguard advantage disappears.

5.6 Experiment 6: Likelihood of renewal

A key assumption of the model is that, due to the dynamic nature of merit, eventual epistemic gain is hard to evaluate. If this assumption is carried over to the science policy domain, it might imply that instead of justifying continued funding based on actual epistemic contributions, funding bodies may be able to spin a tale about the benefits of the research that has been historically funded, regardless of its relative merit in relation to other projects that could have been funded. For a funding body that pursues this 'hype up past results' strategy, a funding mechanism would be preferable not because it leads to high epistemic gains, but

[9]I would like to thank an anonymous reviewer for raising this concern.

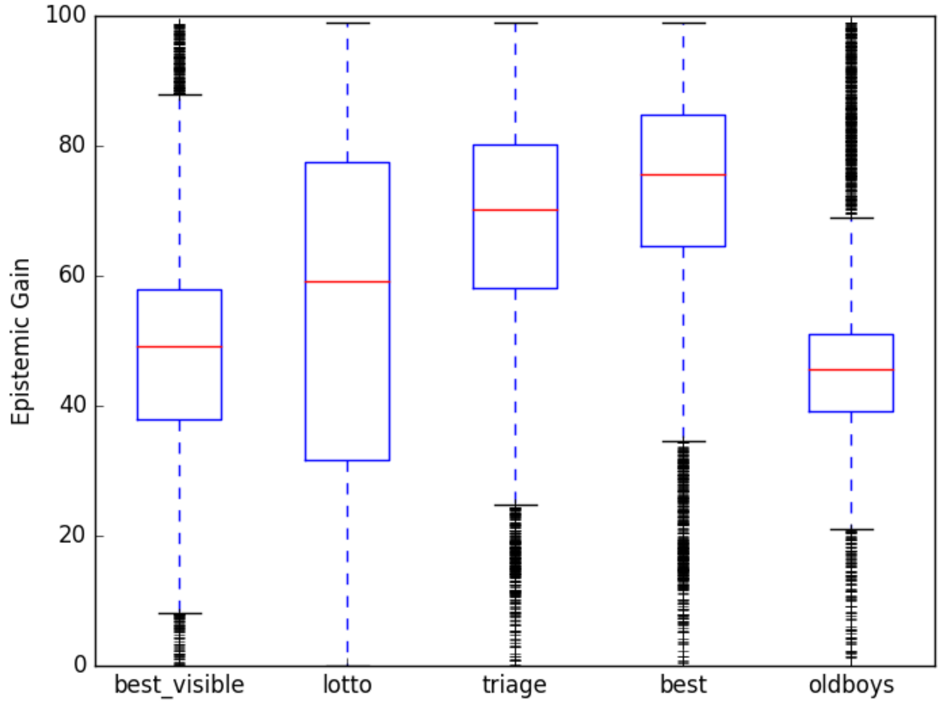

(a) Only 'Winner takes it all'.

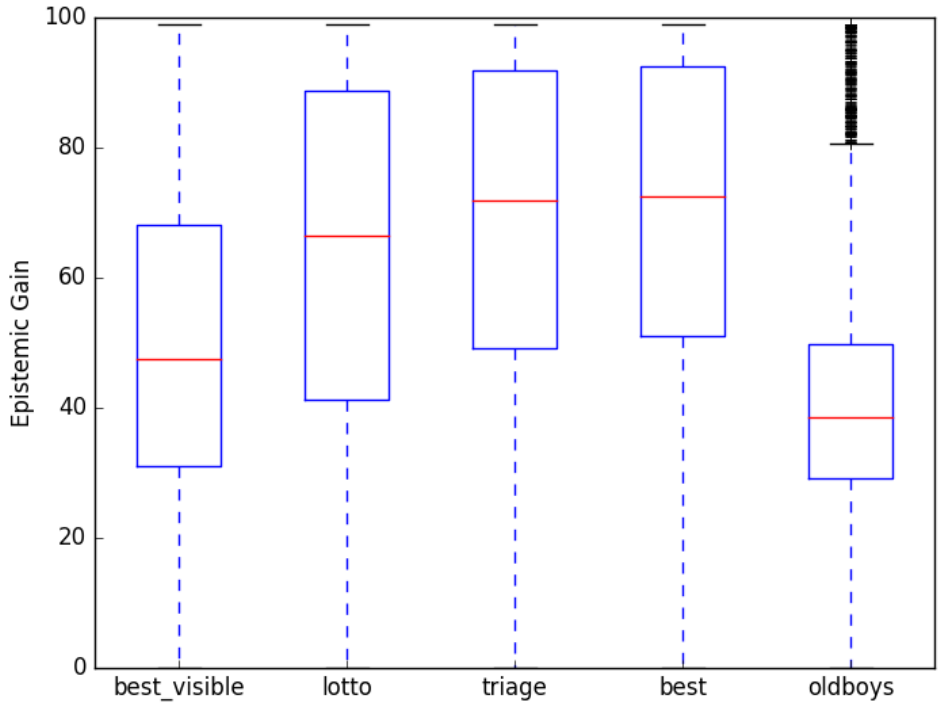

(b) All three dynamic mechanisms.

Figure 7: Comparison of individual epistemic contributions under different funding mechanisms.

because it has a high likelihood of directing funds to previously-funded researchers.[10] The most conservative funding mechanism, 'oldboys', was specifically designed to only provide funds to historically funded individuals (100% renewal). How well do the other funding mechanisms perform in terms of providing fund renewals? To test this, I've compared the renewal likelihood for individuals across the 50 years represented in the simulation, across the different funding mechanisms, using a box-and-whisker plot to demonstrate the distribution of values. Results are presented in Fig. 8 for a (150×150) landscape, in the presence of all three dynamic mechanisms.

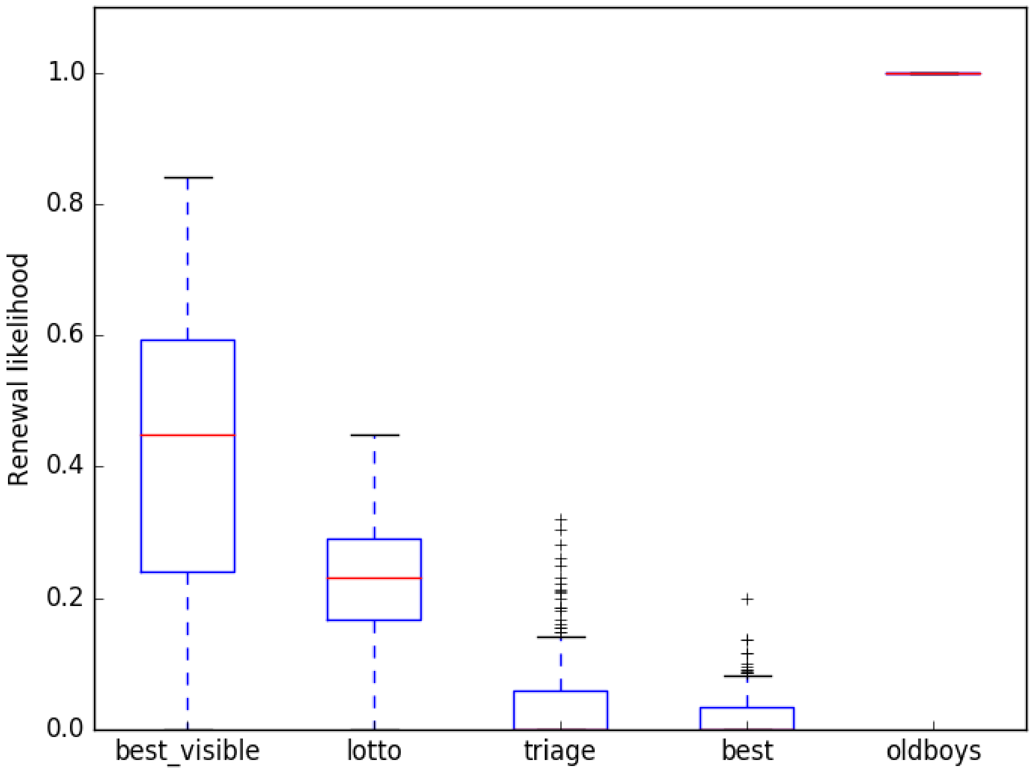

Figure 8: Comparison of renewal likelihood under different funding mechanisms.

[10]There are also less cynical reasons to favour mechanisms that provide continuous funding, such as buildup of tacit knowledge, re-use of existing facilities, and the attraction of stable career paths, none of which are included in the current model and may be explored in future work.

6 Discussion

The key result is that the choice of funding mechanism has a significant effect on the scientific community's ability to generate impactful research, at least for sufficiently large landscapes (size 100×100 and above). Landscape size represents the number of distinct approaches or projects that can be undertaken to explore a scientific topic, with the two landscape dimensions (x and y) representing two independent ways in which projects could differ. A landscape of total size 10,000 represents a situation where we can conceive of 10,000 distinct hypothetical grant proposal all addressing the same topic (though some would be of much greater merit than others). This seems too high if we think of a very specialised topic, but, since the context of this paper is science funding, we should consider all the projects that get lumped together into a single 'study section'. Under this reading, a 'topic' can be as broad as 'search for novel cancer therapies' or 'search for better understanding of bio-molecular cellular mechanisms involved in development', and the number 10,000 seems, if anything, too low, even given the coarse-grained interpretation of the landscape coordinates argued for in §4.3. Furthermore, interdisciplinary links between topics, now commonplace, may generate joint landscapes that are much larger than the specific landscapes corresponding to each disciplinary topic area. Nonetheless, this context sensitivity should be kept in mind both for policy relevance (not all research fields are 'broad' in the relevant sense) and for experimental testing of the results shown above.

Funding mechanisms have a significant effect, and size matters, but which mechanisms perform best? As is clear from the simulations, when only the 'Winner takes it all' dynamic mechanism is in play the 'best' funding mechanism is indeed best at accumulating significance over time, though with various lead margins over the second best strategy. However, as mentioned above, the 'best' funding strategy is not realisable, as it requires a 'god's eye' view of the epistemic landscape. Interestingly, when all three dynamic mechanisms are at play, the performance of 'best' on large landscapes is, within simulation variability, similar to that of 'lotto' and 'triage'. This suggests that the specific ways in which scientific merit is time dependant, a phenomenon which is generally recognised but poorly characterised, has direct relevance to the design of scientific institutions.

An important factor in the epistemic success of peer-review-based funding mechanisms relative to non-peer-review-based mechanisms is the 'vision range' of scientists. Grant peer review, represented in the model by 'best_visible', is considered the 'gold standard' of *ex ante* project evaluation and selection, though in the simulations it often performs worse than 'lotto' and 'triage'. As expected, when the ability of reviewers to accurately assign merit to novel projects increases (greater 'vision rage'), the relative success of 'best_visible' increases relative to 'oldboys', 'lotto', and 'best' ('triage' also benefits from increased vision range). The empirical work of Boudreau *et al.* ([2016]) mentioned in §4.3 suggests that reviewer bias against novel proposals is a real phenomenon. The model presented here sketches how this bias could lead to significant loss of epistemic opportunity on the community level, and why we might want to opt for a funding mechanism that is not (or not entirely) based on peer review.

In addition to relative success with respect to long-term communal epistemic gains, the simulations were also used to evaluate two other metrics that may be important for funding institutions: variability of individual epistemic contributions, and likelihood of renewal of funding. One reason funding bodies might seek low variability is their accountability to the taxpayer, and they may be asked by the taxpayers' representatives to show epistemic progress on an annual level or even on a single grant level (Brennan [2013]). The simulations show that 'best_visible' does result in lower variability compared to 'lotto', though when all three dynamic mechanisms are at play this advantage is small and the epistemic price paid is high. This underscores the above call for further study of merit dynamics, as it may offer an in principle reason for why governance of science at the single project level is likely to be counterproductive.

Since evaluation of individual epistemic contributions is hard, another way to increase accountability is to seek a high renewal rate, with the underlying reasoning that experienced scientists are better placed to undertake new research projects relative to inexperienced scientists. The present model challenges this view by representing a counter argument: new entrants to the field are not encumbered by previous intellectual and material investments, and are therefore more likely to work on novel or neglected topics. Such topics are likely to be of higher epistemic value, especially given merit dynamics. Thus, while ‘oldboys’ and ‘best_visible’ display high renewal rates and relatively low long-term epistemic gains, ‘lotto’, ‘triage’, and ‘best’ display low renewal rates and relatively high long-term epistemic gains. This may be surprising from a policy and intuition perspective, but follows the thrust of previous work in social epistemology, that shows increased diversity and exploration significantly contributes to collective epistemic performance.

The kind of diversity that contributes to communal epistemic success in the current work is cognitive diversity, namely the range of different (and dissimilar) projects pursued by the community. Cognitive diversity is linked to individual diversity (the range of individuals that make up the scientific community) through the introduction of project-choice inertia. An important aspect of diversity highlighted by the model is diversity over time, represented by a low renewal rate, which promotes exploration of novel and neglected areas of the landscape. It should be noted that the model implicitly favours this kind of diversity, by introducing dynamic processes that reduce local merit (through reduced novelty) while increasing merit at potentially distant locations (via new research avenues). These processes are historically and theoretically grounded, but more research is required to evaluate their frequency and magnitude, and to explore other possible mechanisms that may reinforce or contradict their effects. In addition, there are reasons to think that allocation of resources to experienced researchers would have epistemic benefits despite reduced diversity, which would put pressure on the results presented here. Hopefully these will be explored in future models.

7 Conclusion

This paper presents a way to extend existing epistemic landscape models so that they can represent selection by a central funding body. The modified 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 is 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 work 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. The current work looks at the effects of an external, institutional factor, with 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.

Of course, the results from a single idealised model should not be used to directly shape policy. Nonetheless, the current model has a role to play within the growing literature looking at alternative funding mechanisms (Brezis [2007]; Gillies [2008], [2014]; Graves *et al.* [2011]; Herbert *et al.* [2013]; Bollen *et al.* [2014]; Fang and Casadevall [2016]). As current institutions are being questioned, the importance of conceptual clarity on key terms increases. What do we mean by 'merit' and 'impact'? What notions of 'diversity' are useful in this context? Do bibliometric indicators provide relevant evidence for comparing funding mechanisms? These questions are being debated outside philosophy of science, with real policies at stake. In New Zealand, the Health Research Council has started a pilot of funding projects using a gated lottery (Barnett *et al.* [2015]). Meanwhile, funding bodies in Australia are actively opposed to this mechanism (Barnett [2016]). The field is known to be hard to study due to limited access to the grant review process (Chubin and Hackett [1990], [1994]) and due to methodological difficulties in evaluating the impact of research (Dinges [2005]). For formal social epistemology there could be no better testing grounds for the modelling techniques developed over the past decade: can our models fix useful concepts and provide templates for causal mechanisms that could be at play? Could they be used to help shape the debate around emerging policy decisions? The answers will come from future work in the field.

Acknowledgements

The paper is based on PhD research supervised by Tim Lewens and Stephen John, who provided much needed guidance, and examined by Anna Alexandrova and Donald Gillies, who suggested helpful corrections. In bringing this paper to publication I have benefited from the advice of Anna Alexandrova, Steve Wooding and Adrian Currie. Versions of this work have been presented at the conferences of the Society for the Philosophy of Science in Practice, the European Philosophy of Science Association, and the Philosophy of Science Association, as well as at the CamPoS group at the University of Cambridge and at RAND Europe. The paper has been much improved by comments from two anonymous referees.

Shahar Avin

University of Cambridge

Centre for the Study of Existential Risk

16 Mill Lane

Cambridge, UK, CB2 1SB

sa478@cam.ac.uk

References

Alexander, J. M., Himmelreich, J. and Thompson, C. [2015]: 'Epistemic landscapes, optimal search, and the division of cognitive labor', *Philosophy of Science*, **82**(3), pp. 424–453.

Allen, G. E. [1975]: *Life science in the twentieth century*, History of science. New York: Wiley.

Barnett, A. G. [2016]: 'Funding by lottery: political problems and research opportunities', *mBio*, **7**(4), pp. e01369–16.

Barnett, A. G., Graves, N., Clarke, P. and Blakely, T. [2015]: 'What is the impact of research funding on research productivity?', Http://eprints.qut.edu.au/83127/.

Bollen, J., Crandall, D., Junk, D., Ding, Y. and Börner, K. [2014]: 'From funding agencies to scientific agency', *EMBO reports*, **15**(2), pp. 131–133. <http://embor.embopress.org/content/15/2/131>

Boudreau, K. J., Guinan, E. C., Lakhani, K. R. and Riedl, C. [2016]: 'Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science', *Management Science*, **62**(10), pp. 2765–2783.

Brennan, P. [2013]: 'Why I Study Duck Genitalia', http://www.slate.com/articles/health_and_science/science/2013/04/duck_penis_controversy_nsf_is_right_to_fund_basic_research_that_conservatives.html, Accessed 24 August 2014.

Brezis, E. S. [2007]: 'Focal randomisation: An optimal mechanism for the evaluation of R&D projects', *Science and Public Policy*, **34**(10), pp. 691–698.

Brock, W. and Durlauf, S. [1999]: 'A formal model of theory choice in science', *Economic theory*, **14**(1), pp. 113–130.

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

Chubin, D. and Hackett, E. [1990]: *Peerless science: Peer review and US science policy*, Albany: State University of New York Press.

Chubin, D. E. [1994]: 'Grants peer review in theory and practice', *Evaluation review*, **18**(1), pp. 20–30.

De Langhe, R. [2014]: 'A unified model of the division of cognitive labor', *Philosophy of Science*, **81**(3), pp. 444–459.

Dinges, M. [2005]: *The Austrian Science Fund: Ex post evaluation and performance of FWF funded research projects*, Vienna: Institute of Technology and Regional Policy.

Fang, F. and Casadevall, A. [2016]: ‘Research funding: the case for a modified lottery. mBio 7: e00422-16’, .

Gavrilets, S. [2004]: *Fitness landscapes and the origin of species*, Monographs in population biology. Princeton, NJ: Princton University Press.

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).
<http://riviste.unimi.it/index.php/roars/article/view/3834>

Goldman, A. and Blanchard, T. [2016]: ‘Social Epistemology’, in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, winter 2016 edition.

Graves, N., Barnett, A. G. and Clarke, P. [2011]: ‘Funding grant proposals for scientific research: retrospective analysis of scores by members of grant review panel’, *BMJ*, **343**.

Greenberg, D. S. [1998]: ‘Chance and grants’, *The Lancet*, **351**(9103), pp. 686.

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), pp. 363–374.

Herbert, D. L., Barnett, A. G., Clarke, P. and Graves, N. [2013]: ‘On the time spent preparing grant proposals: an observational study of Australian researchers’, *BMJ Open*, **3**(5).

Kennedy, J. V. [2012]: ‘The Sources and Uses of U.S. Science Funding’, *The New Atlantis*, **36**, pp. 3–22.
<http://www.thenewatlantis.com/publications/the-sources-and-uses-of-us-science-funding>

Kitcher, P. [1990]: ‘The Division of Cognitive Labor’, *The journal of philosophy*, **87**(1), pp. pp. 5–22.
<http://www.jstor.org/stable/2026796>

Kitcher, P. [1993]: *The advancement of science*, New York: Oxford University Press.

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

Kuhn, T. [1970]: *The structure of scientific revolutions*, Chicago: University of Chicago Press.

Lee, C. J. [2015]: 'Commensuration Bias in Peer Review', *Philosophy of Science*, **82**(5), pp. 1272–1283.

Muldoon, R. and Weisberg, M. [2011]: 'Robustness and idealization in models of cognitive labor', *Synthese*, **183**(2), pp. 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., Elliott, K. C., Haufe, C. and Burian, R. M. [2009]: 'Philosophies of funding', *Cell*, **138**(4), pp. 611–615.

Polanyi, M. [1962]: 'The Republic of Science: Its political and economic theory', *Minerva*, **1**, pp. 54–73.

Pöyhönen, S. [2016]: 'Value of cognitive diversity in science', *Synthese*, pp. 1–22.

Stirling, A. [2007]: 'A general framework for analysing diversity in science, technology and society', *Journal of the Royal Society Interface*, **4**(15), pp. 707–719.

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

Thoma, J. [2015]: 'The Epistemic Division of Labor Revisited', *Philosophy of Science*, **82**(3), pp. 454–472.

Weisberg, M. and Muldoon, R. [2009]: 'Epistemic Landscapes and the Division of Cognitive Labor', *Philosophy of science*, **76**(2), pp. 225–252. <http://www.jstor.org/stable/10.1086/644786>

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

Zollman, K. J. [2010]: ‘The epistemic benefit of transient diversity’, *Erkenntnis*, **72**(1), pp. 17–35.