

Explorers, Extractors, and the Benefits of Diversity in Science

1. Introduction

Scientists differ in the ways they approach their work. Some are happy to follow in the footsteps of others, and continue with work that has proved to be fruitful in the past. Others like to explore novel approaches. It is tempting to think that herein lies a division of labour that is conducive to overall scientific progress: The latter, explorer-type scientists, point the way to fruitful areas of research, and the former, extractor-type scientists, more fully explore those areas. Still, showing that such a division of labour is indeed beneficial with a formal model has not proven easy. While Weisberg and Muldoon (2009) show that a group of scientists that use what they call a ‘follower’ strategy do better when there are also ‘mavericks’ around, their results suggest that it would be better still if all scientists were mavericks. Their model thus does not provide a full-fledged demonstration of the benefits of division of labour. I here present a model which can show the benefits of such cooperation. However, my model also suggests that this division of labour is unlikely to emerge without special incentives, since within the cooperative schemes, explorer-types always do worse at finding significant scientific results than extractor-types. I suggest that a bonus in personal reward for novel research results could play such a role, and that thus the value we tend to place on novelty is an incentive mechanism to sustain a beneficial division of labour.

I will study the division of labour between explorer-type scientists who seek novel discoveries, and more conservative extractor-type scientists when exploring a field of research, in which there are many different discoveries to be made - some smaller, some larger. Such a research field can be studied well with an agent-based ‘epistemic landscape’ model as Weisberg and Muldoon (2009) introduce it, where scientists move around a previously unknown landscape, making subsequent discoveries. While Weisberg and Muldoon show that scientists who take into account what others have done do better than scientists who who work in apparent isolation (which already amounts to a kind of division of labour), in their model, pure groups of ‘mavericks’ always do best. One intuitive explanation of why this is so is that their ‘follower’ scientists end up duplicating the work of others much of the time: They are characterised by a tendency to move to the most successful previously explored approaches in their neighborhood. But mere

duplication has no added benefit in the setup of the model - Weisberg and Muldoon's measure of success only counts the first time an approach is used¹. But the frequency of mere duplication, I argue, is also a feature of the model that is not a very plausible description of extractor-type behaviour in science.

My model studies the behaviour of scientists who all avoid merely duplicating the work of others. In that sense, all the agents in my model behave like Weisberg and Muldoon's 'mavericks', who always go for a previously uninvestigated approach when possible. I think that given the modelling assumption that all scientists successfully uncover all significant results from using an approach the first time they use it, this makes sense for two main reasons: First, research results, especially those of researchers working on approaches similar to one's own, are mostly freely available, so we can assume that researchers already have access at least to the research results from approaches that are similar enough to their own for them to understand and easily find them. Second, if this is so, it is hard to see why anybody could be motivated to simply duplicate the work of others. There is no epistemic benefit from doing so, nor do there seem to be any social rewards from mere duplication. So a model which has all agents avoid duplicating the work of others gives a more plausible description of scientists' behaviour. And it also carries the hope of showing division of labour to be beneficial.

We can still meaningfully distinguish between extractor-type and explorer-type scientists even when all avoid duplication, in the following way. My model describes Explorers as scientists who like to follow approaches that are very different from those of others, while Extractors like to do work that is very similar to but not the same as that done by others. Given these two types of strategies, my central question is: Are mixed populations of Extractors and Explorers better at making scientific discoveries than pure populations of Extractors and Explorers respectively? The results I will present suggest that the answer to this question depends on how large the range of movement and awareness of other scientists' work is: Division of labour is more beneficial when this range is large. But the only case in which division of labour is not beneficial is that of local movement, which is also the only kind of movement Weisberg and Muldoon study, and which, I will argue, is less relevant when we think of actual scientific progress.

This paper will proceed as follows: Section 2 introduces my epistemic landscape model. Sections 3 and 4 present the results from simulations run on the model for the case of local movement, and medium range

¹ Their primary measure of success is 'epistemic progress', that is, the proportion of significant approaches that have been discovered, i.e. used at least once. Hence the subsequent uses of previously discovered approaches do not count. Weisberg and Muldoon suggest that division of labour may end up being optimal after all when we take into account that the 'maverick' strategy is more expensive than the 'follower' strategy. Still, without a formal model, this is not yet convincing: After all, even if it is cheaper, following may not be worth its money if it produces no added value much of the time.

as well as global movement respectively. Section 5 suggests some explanations of why the model behaves as it does, and section 6 discusses the relevance of these results for real scientific contexts.

2. The Model

The model I implemented consists of an epistemic landscape, and scientist-agents who move on that landscape. An epistemic landscape is meant to represent a single research topic. Different points in the epistemic landscape represent scientific *approaches* with an associated epistemic *significance*. An approach is characterised by a methodology (regarding data-collection as well as data analysis), a research question and set of background beliefs. Epistemic landscapes represent similarity of research approaches as spatial proximity. Like Weisberg and Muldoon's, my epistemic landscape model is two-dimensional and discrete, so that different research approaches are represented as discrete patches on a plane, arranged on a grid. These research approaches can vary along two dimensions, namely the x- and y-axis. Even though research approaches are assumed to be discrete, we can interpret them to be very fine-grained: For instance, using the same methodology and background assumptions as somebody else in order to treat a slightly different research question could count as a different, but nearby approach. The epistemic significance associated with each approach is a numerical value which represents, roughly, the amount of scientifically or socially important results that can be obtained using a particular approach². My model, like Weisberg and Muldoon's, assumes that epistemic significance is not distributed randomly on the landscape, but in two Gaussian-shaped hills with single peaks, where one of the peaks is higher than the other. This makes sense when we assume that approaches that are similar to significant approaches are also likely to be similarly significant, and that there are several areas of epistemic significance within one research field.. Large parts of the epistemic landscape are assumed to be entirely insignificant. I study an epistemic landscape with a 101 x 101 grid. Fig. 1 shows an epistemic landscape where the patches are coloured according to significance.

² See Kitcher (1993) on the concept of epistemic significance. Weisberg and Muldoon stress that epistemic landscape models do not rely on any particular interpretation of epistemic significance, but do assume that we all share the same conception of epistemic significance, and correctly identify it. This is because all agents make their behaviour depend on one single value of epistemic significance for each approach, which is the same for all. This is, of course, an idealisation, which could be relaxed in further studies.




Fig. 1: An Epistemic Landscape

Scientific progress is now modelled as occurring when scientist-agents move around the landscape making discoveries. When scientists ‘visit’ a patch, they use the approach of that patch to find out its significance. Again like Weisberg and Muldoon, I assume that all agents successfully determine the significance of an approach when they use it³. Scientists can move around the landscape according to rules. By specifying different rules, we can study how well different types of scientists do at exploring the epistemic landscape, and how they interact, in the hope of thereby learning something about scientific progress in the real world. I want to use this epistemic landscape model to understand the interaction of explorer-type and extractor-type scientists, that is, scientists who like to do things that are maximally different from what others are doing, and scientists who like to do similar work to others, respectively. Due to its spatial representation of similarity, the epistemic landscape model lends itself very well to implementing and studying the effects of these types of strategies.

Scientist-agents are characterised by their position on the landscape, their heading, their memory of what approaches they have used, and a rule for moving around the landscape, which is responsive to what approaches have been used, and to what success, in the area around them. Focusing first on the case where agents move only to patches in their neighborhood, and respond only to what has happened in their neighborhood, the rules I will investigate can be described as follows in pseudo-code:

³ Relaxing this assumption is another way in which one could make Weisberg and Muldoon’s ‘follower’ strategy more plausible: If the first person to use an approach didn’t uncover all significant results that can be obtained with that approach, there would again be a point to ‘follower’ behaviour. This would surely be another interesting area for investigation. I choose the option of changing the strategy to avoid duplication in order to keep as much as possible of the original framework fixed and for ease of presentation.

Explorer Rule:

Ask: Have any approaches in my Moore neighborhood been investigated by other scientists?

If no: Ask: Does my current approach have equal or greater significance than my previous approach?

If yes: Move one forward.

If no: Go back one patch and set a random new heading.

If yes: Ask: Are there any approaches in my neighborhood which have not yet been visited?

If yes: Move to the unvisited neighboring patch at the greatest minimum distance to my neighbouring patches previously visited by other scientists. If several patches are at an equal distance, pick randomly between them.

If no: Pick a random new approach in the neighborhood.

This rule expresses both a desire to make significant discoveries, as well as a desire to get further away from other scientists. If they are in previously uncharted territory, Explorers behave like ‘hill-climbers’, and change course whenever they start going downhill. If they encounter previously visited approaches, they go in the direction that faces away from these approaches (to the patch at a maxi-min distance to the visited patches), and continue going that direction unless they start going downhill or encounter other previously discovered patches. In the context of the epistemic landscape, this strategy captures the mentality of a scientist who wants to make discoveries that are as different from those of others as possible.

Extractor Rule:

Ask: Are there any patches in my Moore neighborhood that have been visited by other scientists?

If no: Ask: Are there any unvisited neighboring patches?

If yes: Pick a random one of these patches.

If no: Pick a random new patch.

If yes: Ask: Are there any unvisited neighboring patches?

If yes: Go to the unvisited patch in my neighborhood that is closest to the approach in my neighborhood previously investigated by other scientists with the highest significance. If there are several such approaches, pick randomly between them.

If no: Go to a random neighboring approach.

This rule expresses both a desire to make significant discoveries (since there are likely to be significant discoveries to be made in the neighborhood of significant discoveries), as well as a desire to do things that are similar to what others are doing. Still, extractors, too, will go for an uninvestigated approach whenever there is one in the neighborhood - they do not want to merely duplicate the work of others. This hence captures the mentality of scientists who, while concerned to make significant discoveries, want to stay close to what other scientists are doing.

I assume that scientist-agents are initially distributed randomly across insignificant areas of the epistemic landscape, and then start moving around the landscape according to the above rules. Each agent applies their rule and moves once at each time step, or round. To study their behaviour, I ran simulations where I let different groups of scientist-agents move and interact for a fixed number of rounds, and measured how well they did. But what would it mean for a group of scientists to be successful at exploring an epistemic landscape? What matters, and what I take to be the most important measure of success, is what proportion of the total significance the community of scientists discovers, and how fast it does so - where the total significance is the sum of the epistemic significance of each of the individual approaches, that is, all the significance to be discovered in the research field. My discussion will hence focus on this measure of success. In addition, I collected data on the average significance of the moves made by the different types of scientists, in order to compare how well each of the groups does in various scenarios.

Before presenting the results from the simulations I ran, let me stress that the above model is still highly idealised in many ways. For instance, real research fields will be comprised of approaches that vary along more than two dimensions. A more detailed model would also incorporate costs of movement and research, and differences in talent and speed. It may contain various different kinds of distributions of epistemic significance, for instance ones with more than two hills. Section 6 will provide a discussion of what I believe we can nevertheless learn from this model.

3. Results: Local Movement

I first ran a number of simulations in order to see how well pure groups of Explorers and Extractors do at exploring the epistemic landscape. For each type of scientist, I looked at the behaviour of groups of 20, 40, 60, 80 and 100 agents over the first 500 rounds, and repeated simulations for each of these model specifications 50 times. Fig 2 shows the proportion of significance discovered, on average, by the groups of scientists relative to the number of rounds that have passed. The graphs show that for both Explorers and Extractors it is true that the larger the group, the faster they find significant approaches. It also appears that both Explorers and Extractors eventually discover all, or nearly all significant approaches:

All groups apart from the 20 agent groups on average discovered more than 97% of the total significance after 500 rounds (with a very small standard deviation in all cases).




Fig. 2: Average proportion of total significance discovered by groups of different sizes

Next, when comparing how well pure groups of Extractors do relative to Explorer groups of the same size, it turns out that Explorers always do better than Extractors. Fig. 3 shows this for the case of 20 and 40 agents.




Fig. 3: Proportion of Total Significance: Explorers and Extractors compared

Another striking difference between the pure groups of Explorers and Extractors is that the standard deviation for the Extractors tends to be higher than that for the Explorers, especially when the group size is small - there is hence more variation in how well the groups of Extractors do. Looking at how the

scientist-agents move on the epistemic landscape may illuminate why these differences occur. Fig. 4 shows the movements of typical groups of 40 Extractors and 40 Explorers respectively up to round 300 (all scientist-agents left coloured traces of their movements). The Extractors here still haven't found the most significant area of one of the hills, probably because of their unwillingness to stray away from the previously discovered areas, which, in the beginning, are all in the insignificant area of the landscape. The hill-climbing activity of the Explorers, on the other hand, seems to have lead them to the most significant areas of the landscape relatively quickly, and does so in a similar way in every simulation.




Fig. 4: 40 Explorers (left) and 40 Extractors (right) moving on the epistemic landscape

In order to study the behaviour of mixed groups of Explorers and Extractors, I focused on the fixed group size of 40 agents. In addition to the pure groups, I looked at mixes with 10%, 20%, 30%, etc. Explorers, again running simulations for each model specification 50 times. Fig. 5 illustrates how much of the total significance has been discovered, on average, by each of the groups after 50, 150 and 300 rounds, and shows a more detailed time evolution for the pure groups as well as a 50/50 mix of Explorers and Extractors below.

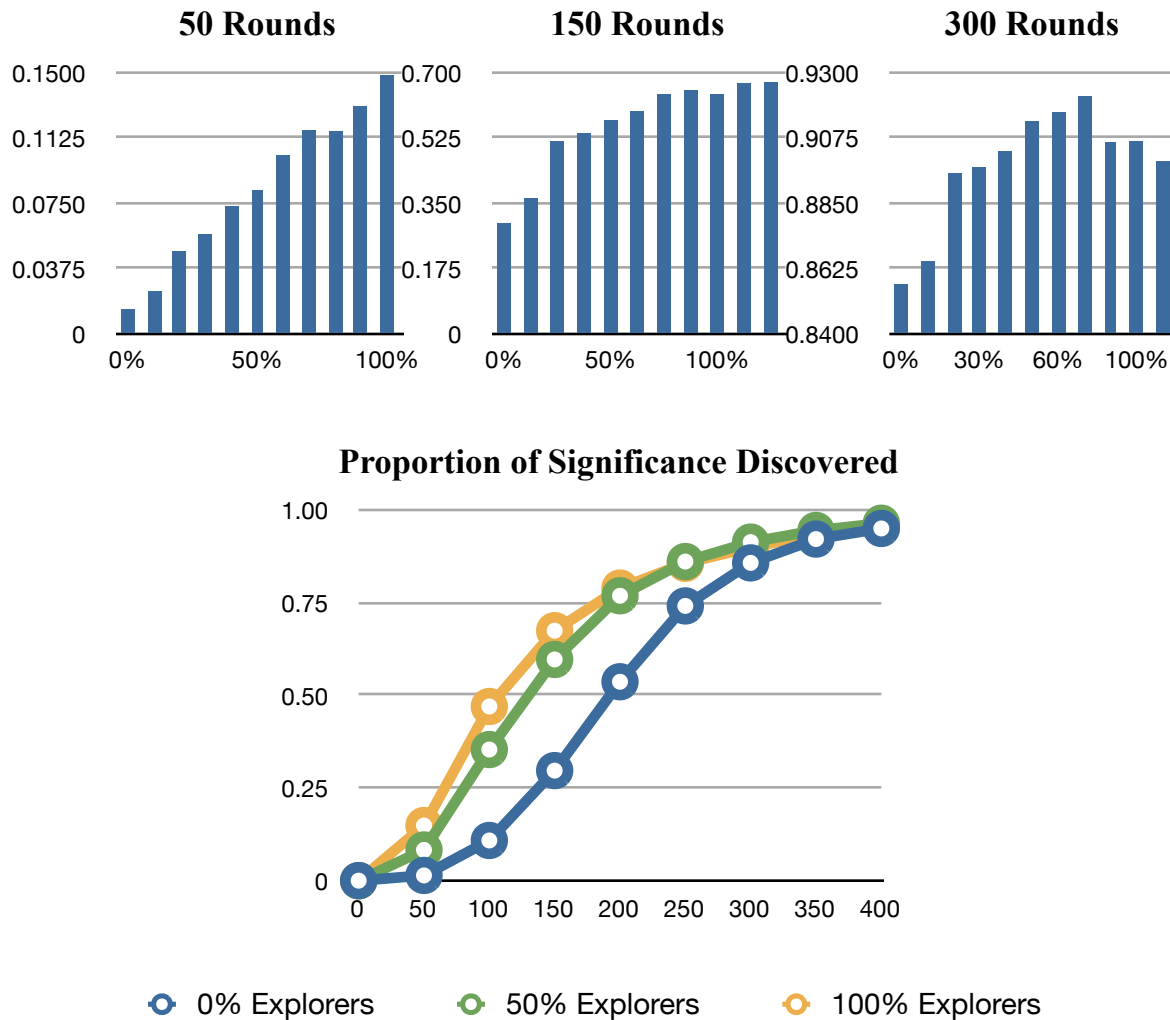


Fig. 5: Mixed groups (local movement)

What emerges here is that until fairly late, pure Explorer groups are faster than mixed groups at discovering the significant approaches in the landscape. After 50 rounds, Explorers have discovered 14.9% of the epistemic significance on average, while groups with an equal mix only discovered 8.3% and groups of Extractors only 1.4%. After 150 rounds, there is still a significant difference. Here the groups have discovered 67.4%, 59.7%, and 29.7% of the landscape on average, respectively. After that, however, the mixed groups catch up and do similarly well as the Explorers. At 250 rounds, they have found roughly equal amounts of significance, and at 300 rounds, the mixed groups have done slightly better than the Explorers: The groups with an equal mix then discovered 91.3%, whereas the Explorer groups discovered 89.9%. The group that has done the best after 300 rounds is in fact the mixed group with a mix of 70% Explorers and 30% Extractors, which has discovered 92.2% on average. Still, the differences between the Explorer group and the mixed groups are very slight, and emerge only when a large part of the total significance has already been discovered. Early on, Explorers certainly discover the significant parts of the landscape much faster than the mixed groups. In most contexts, where we are

interested in early and medium run success, this should make groups of pure Explorers the most successful at realising the goal of finding significant approaches fast.

Note, however, that so far, we have only looked at the case where agents are aware of only what happens in their immediate neighborhood, and can only move to a neighboring patch. This is what I want to call ‘local movement’. We have seen that the extent to which division of labour is beneficial when movement is local is very limited. As the next section will show, things are different once agents have a slightly larger range of movement and awareness.

4. Results: Medium Range and Global Movement

We have seen that when movement is local, and unless we care most about how fast we approach discovering the total significance of the landscape in the long run, it still pays off to have pure groups of Explorers. Interestingly, however, division of labour becomes much more beneficial when we increase the range within which scientist-agents can move on the epistemic landscape, and within which they are aware of previous discoveries. To study movement in larger ranges, I adapted the Explorer and Extractor Rules by simply replacing ‘Moore neighbourhood’ by the radius the agents consider. For instance, in the case of the Extractors, if possible, agents move to the patch within a radius x which is closest to the best previous discovery within the same radius x .

Movement within a range of 1 includes consideration of only four neighbors, and hence less than the Moore neighbourhood, and is hence even more restricted than the local movement I considered above. However, movement within a range of 2 or more means that the scientist-agents have an expanded horizon: they are more flexible with respect to what new approaches they can adopt, and their choices respond to previous discoveries made using approaches that are more different from their own. In order to study whether division of labour between Explorers and Extractors is beneficial when agents are more flexible in this sense, I ran simulations with agents moving within ranges 1-5, 7 and 10. As before, I restricted my analysis to groups of size 40, and looked at group compositions with 0%, 10%, 20% and so on of Explorers.



As might be expected, at range 1, we get similar results as with local movement. In fact, the pure group of Explorers does best here even after 300 rounds. However, as the range gets larger, mixed populations do better and better compared to the pure populations. At range 2, the pure Explorer-group still does better after 50 time steps, but after time step 100 already, the mixed group with 70% Explorers did best. At range 3 and larger, there were mixed populations that did better on average than both pure populations for

every time step I considered. Fig. 6 gives an overview of which mix did best after all the time steps, along with the proportion of the total significance that mix discovered on average by that time step.

#Rounds	50		100		150		200		250		300	
	Ideal %Explorers	%Signif.	Ideal %E	%S	Ideal %E	%S	Ideal %E	%S	Ideal %E	%S	Ideal %E	%S
Range 1	100%	13.70%	100%	45.60%	100%	66.80%	100%	78%	100%	84.80%	100%	88.80%
Local	100%	14.90%	100%	47.00%	100%	67.40%	60%	79.20%	70%	87.30%	70%	92.20%
Range 2	100%	13.50%	70%	46.30%	70%	73.50%	30%	88.20%	30%	95.20%	30%	98.40%
Range 3	80%	15.20%	50%	52.40%	30%	82.90%	30%	94.80%	20%	98.90%	20%	99.80%
Range 4	50%	17.60%	30%	60.30%	30%	86.90%	30%	95.60%	30%	97.90%	30%	98.20%
Range 5	30%	21.30%	30%	64.40%	20%	89.20%	20%	97.70%	20%	99.70%	60%	100.00%
Range 7	70%	31.70%	40%	72.90%	40%	93.20%	40%	98.90%	40%	99.90%	70%	100.00%
Range 10	50%	43.20%	30%	77.40%	30%	93.50%	30%	98.50%	50%	99.80%	80%	100.00%
Global	50%	35.90%	40%	53.70%	70%	62.50%	90%	69.40%	90%	74.50%	90%	78.70%

Fig. 6: Optimal group compositions at different ranges of movement

The exact numbers should be treated with caution here, however. The differences between the different mixed groups are often very small, and the standard deviations for the groups that involve more Extractors, and for the smaller ranges, tend to be quite large (up to 15 percentage points). Still, the differences between the pure populations and the mixed populations that did best are large enough for many ranges for us to be fairly confident that division of labour is beneficial for these ranges of movement. Out of the ranges I considered, the case of range 10 is most pronounced. Fig. 7 illustrates the results for this range. The mean values for the proportion of significance discovered by the different groups after 50, 100 and 150 rounds is clearly highest for mixed populations: The top graphs are hill-shaped, with a relatively flat top and steep edges. Below is a more detailed comparison of the pure groups and the mixed group with 40% Explorers. The difference here between the mixed group and the pure Explorer group is fairly large, and for both groups the standard deviation is small. The difference between the mixed group and the pure group of Extractors is less big, while the standard deviation for the pure group of Extractors is quite large, especially at earlier time steps. So while the advantage of the mixed group over the Extractor group in terms of average proportion of significance discovered is less big, the mixed groups are here also preferable for the more stable performance they offer.



	#Rounds	0	50	100	150	200	250	300
0% Explorers	Mean	0	0.2487358	0.6387783	0.8535753	0.9438267	0.9778898	0.9903267
	Stand. Deviation	0	0.0825226	0.1204479	0.1249546	0.0843507	0.0536223	0.0459869
40% Explorers	Mean	0	0.4229383	0.772375	0.933228	0.9835107	0.9957914	0.9997326
	Stand. Deviation	0	0.0475176	0.0385805	0.0309568	0.0262377	0.0185531	0.0015897
100% Explorers	Mean	0	0.1971685	0.3905726	0.5843849	0.7804787	0.9666811	1
	Stand. Deviation	0	0.0171182	0.0211271	0.0242372	0.0179282	0.011808	1.994E-15

Fig. 7: Division of Labour at Range 10

Compared to our analysis of the case of local movement, what is also interesting is that at range 10, Extractors do better than Explorers, rather than the other way around. As can be seen in Fig.7, pure groups of Extractors do better at every time step than pure groups of Explorers. And it can in fact be shown that Extractors also do better than Explorers within the mixed groups. To show this, I also recorded data on

how much significance, on average, each type of scientist newly discovers each round. To measure how well each group is doing relative to the other, I assigned the significance of an approach to the group that first discovers it, and divided the total for each group by the number of moves that have been made by the members of that group. Fig. 8 shows that Extractors do better than Explorers according to this measure at each time step when there is a 40% proportion of Explorers in the mix. And indeed, at range 10, Extractors do better for all the time periods and all the group compositions I recorded.




Fig. 8: Average Significance of Moves with 40% Explorers

What the coarseness of the data I recorded hides, however, is that for most simulations I ran, there was a brief time period when Explorers did better than Extractors. This period is more prolonged when the range of movement is smaller. For instance, at range 3, Explorers still do better after 50 time steps, which Fig. 8 also shows. At range 2, Extractors only start doing better than Explorers on average after about 200 time steps. When movement is local, and the mixed groups do worse than pure groups of Explorers, Explorers also do better within the mixed groups.

When looking at ranges smaller than 10, we find that the qualitative results I just described are the same but less pronounced for ranges 3 and up. Range 2 appears to be a kind of boundary case between the local movement we looked at before and this kind of medium range movement: Here Explorers still do mostly better than Extractors, and the differences between the performance of the pure group of Explorers and the mixed groups is still quite small.

So it appears that division of labour becomes more and more beneficial, the larger the range at which agents can move. Moreover, the data also suggests that all types of groups do better the larger the range of

movement (see, for instance, Fig. 6). Given this tendency, studying the case of global movement is also instructive. The Extractor rule can be adapted for the global case simply by having agents move to the uninvestigated patch that is closest to the best previous discovery in the entire research field. For computational reasons, however, I had to adapt the Explorer rule as follows for the global case: As long as there is no patch that has previously been visited by another agent in a radius of 10 patches, global Explorers behave like hill-climbers. When there is such a patch, they try to find an unvisited patch that has no other patch in their radius of 10 somewhere in the research field. If they don't find one after considering 100 unvisited patches, they go to a random unvisited patch. Although this global rule departs from the Explorer rule we used before, some pilot simulations suggested that this global Explorer behaviour is not significantly different from that of Explorers with range of movement of 71 (which covers the whole research field). After a few rounds, both groups appear to jump randomly between the unvisited patches of the research field.

As Fig. 6 already indicates, division of labour is still optimal in this global case, at least within the time periods I looked at. However, two things change: Firstly, pure groups of Extractors start doing worse than Explorers again after some time, as Fig.9 shows, and the mixes that are best later on are the ones with a high number of Explorers. So there is something about the global context that makes the Extractor strategy perform worse again. Secondly, all groups do significantly worse than the groups that moved at a medium range. While at range 10 all the groups had on average discovered almost the entire total significance after 300 rounds, now the best mix with 90% Explorers discovered 78.7% on average, and the worst group, the pure group of Extractors, discovered only 55.4%.




Fig. 9: Global Movement

So the results suggested by the simulations I ran can be summarised as follows:

Local Movement:

All groups tend to discover all or nearly all of the total significance in the long run. Still, Explorers do better than both Extractors and mixed groups of Explorers and Extractors at finding significant approaches in the earlier stages. Within the mixed groups, Explorers make more significant discoveries than Extractors do.

Medium Range Movement:

At ranges larger than 2, mixed groups do better at discovering significant approaches than pure groups, and thus division of labour becomes beneficial. Extractors also start doing better than Explorers, although the performance of pure groups of Extractors shows a much greater variance. Mixed groups have a smaller variance, do better on average, and within them, Extractors still do better than Explorers (as measured by the average significance of their moves). These results become more pronounced for the larger ranges I considered. All types of groups tend to do better and find all or nearly all significant approaches more quickly when movement is medium range as opposed to local.

Global Movement:

Mixed groups still do best for the time periods I considered, but all groups are much slower at finding significant approaches than when they move at medium range, and Extractors start doing worse than Explorers after a short while.

5. Explaining the Results

Whether division of labour is beneficial in our model hence depends on the range at which the agents can move, and within which they respond to previous discoveries. To understand why, it is again instructive to look at representations of the paths the different scientists took in the different scenarios. First, consider Fig. 10, which shows the epistemic landscape after about 100 rounds both for a typical mixed group of scientists (with 40% Explorers) moving at range 1, and one moving at range 10. In the range 10 group, the Extractors seem to have moved to both peaks relatively quickly, and continue covering both hills from the middle to the edges. This activity keeps the Explorers away from the areas of most significance. What could then explain the better performance of mixed groups of Extractors and Explorers may be that Explorers can point the way to the most significant areas more quickly than Extractors would find them

on their own, but that the Extractors then continue to try out all the significant approaches in that area, while Explorers on their own would have some tendency to go on to less explored areas. This also explains why, within the mixed groups, Explorers make the more significant discoveries for a short time in the beginning, before the Extractors take over. This seems to match quite well the way one may intuitively think division of labour between extractor- and explorer-type scientists may work. But why doesn't the same hold for the case of local movement?




Fig. 10: Range 1 (left) vs. Range 10 (right): 40% Explorers (red) and 60% Extractors (blue) after ca. 100 rounds

What is striking in the left image in Fig. 10 is that very few Extractors have actually made it to the most significant areas of the landscape. In this particular case, one of the hills has been climbed only by Explorers. Why haven't the Extractors followed the Explorers to the most significant areas, as they seem to do in the case of larger ranges? The explanation that suggests itself when looking at how the scientists move on the landscape is that Explorers often block the way for Extractors, and thus keep them from following them, or other Explorers to the most significant areas. In the image below, most of the Extractors seem to be 'fenced in'. Explorers often go straight ahead for a long time. And when movement is local, Extractors often cannot cross the trail of previously explored patches the Explorers left behind, and the fencing in we see in the image occurs. And then it is the Explorers who tend to make the more significant discoveries than the Extractors do. When movement occurs at a larger range, this cannot occur so easily, because Extractors can 'jump over' the trails of previous discoveries Explorers leave behind if it gets them to more significant areas.

Looking at how the discoveries of mixed groups unfold when movement is global similarly suggests an explanation of why the global case is different from the medium range cases in the ways we described. Fig. 11 depicts a typical mixed group moving globally. What happened here is that all the Extractors first move to the peak of the one hill and start using the approaches in the vicinity of that peak. They then all switch to the second hill and do the same, leaving the second hill alone till the end of the simulation - thereby not discovering the significant areas still left around the first hill. The Explorers, on the other hand, appear to move to unvisited approaches on the landscape randomly - so they only discover those significant areas rather slowly. The fact that Extractors completely abandon one of the hills then provides a plausible explanation of why the mixed groups do much worse when movement is global, and why Extractors don't do as well within those mixed groups. In fact, often the Extractors in mixed groups only ever go to the hill with the higher peak. The pure groups of Extractors never discover more than one peak, where this is often the lower of the two peaks.




Fig. 11: Global Movement (40% mix after ca. 50 rounds)

The explanation for this behaviour is relatively straightforward: Extractors go to the unvisited patch that is closest to the best discovery that has been made so far. Once an approach is used that is more significant than the peak of the lower hill, Extractors will move to patches as close as possible to that approach and hence start circling around the higher of the two hills. Eventually, they will also cover the smaller hill, but only in virtue of its being close to the higher of the two peaks. The restricted horizon of groups with a medium range of movement keeps them from completely abandoning one of the hills in this way. Hence when there are several hills in the landscape (and one could imagine many more than we considered here), there is a benefit to movement at a restricted range.

6. Discussion

In the model I presented, division of labour between Extractors and Explorers is beneficial under many circumstances. Does this mean that we have shown division of labour between explorer-type and extractor-type scientists in real scientific contexts to be beneficial? One problem with making this inference is the high degree of idealisation of the model. Can we at least learn something qualitative from this model anyway? One good sign is that at least in the medium range case, Explorers and Extractors interact in a way that seems like a very plausible description of interactions between explorer-type and extractor-type scientists in real scientific contexts: Explorers point the way, find the most significant areas, and Extractors uncover all the significant results in the vicinity. This should give us some confidence that our model captures something right. What we observed in the global case, too, seems like something that could plausibly happen: Explorer-types all work on approaches similar to what is acknowledged to have been the most significant so far, whereas smaller significant areas of research become neglected. In addition, my model suggests that there is more of a danger of this occurring when scientists are too flexible in what research approaches they can adopt, and are aware of the work done in too large an area in the research field. To some extent, then, the model shows that more information and more flexibility could become a bad thing.

But do real scientists behave like the medium (or global) range movers in our model, rather than the local movers? If they were more like the local movers, again we wouldn't have shown division of labour to be beneficial. I think there are two reasons why the medium range case is the most relevant for actual scientific contexts. Remember that closeness in the epistemic landscape is interpreted as similarity of research approach. When movement is local, scientists are only aware of those distinct approaches that are most similar to their own, and they can only adopt a research approach that is equally similar to their own. A medium range means that agents have a larger range of awareness and movement. What would justify us to use a local movement model? It would need to be the case that information about more dissimilar approaches is hard to get by, and that it is very difficult to adopt a new approach that is more dissimilar from one's own. It is surely the case that, even though much information is now freely accessible, acquiring detailed knowledge of areas of research very different from one's own is hard. And it is also hard to acquire the new research tools needed to work on a very different approach. But for local movement to be a plausible description of actual scientific progress, it would need to be the case that this is true for all approaches apart from those that are immediately adjacent to one's own.

Whether local movement is plausible depends to some extent on how fine-grained we interpret research approaches to be. If they are coarse-grained, it is more plausible than if the distinctions between neighboring approaches are very small. As I already suggested above, given the way we implemented the

model, we need to in fact think of the epistemic landscape as very fine-grained. We said that a single scientific agent is enough to discover the entire significance of one approach in one round. If approaches were more coarsely grained, this would be unrealistic: It would then have been better to allow for a patch to be visited multiple times, or for single agents to stay on a patch for longer. But in our model, agents move about the epistemic landscape quite fast. Of course, we could think of each round as lasting a long time. But what we aimed to capture is a more detailed look at the dynamics of research groups, that tracks smaller changes in research approach. Imagine we interpret each round as lasting about 1 month. In that case, the 300 rounds we looked at would represent 25 years of research, which seems like a plausible lifetime for a specialised field of research. But given that interpretation, local movement would imply that scientists are not aware of research that they would be able to conduct themselves within 2 months of research. Similarly, they could not now learn the research tools for using an approach that they will be able to use after just one month of research on something else. All in all, I think that given the fine-grained interpretation of research approaches that is implicit in the model, but still acknowledging some limitations to awareness and research flexibility, the case of medium range movement is most relevant for actual scientific research.

Another reason why the case of local movement may be less relevant in applications is that what keeps the scientist-agents from benefiting from division of labour may be an artifact of the two-dimensional nature of the model. The kind of ‘fencing in’ that we observed is on the one hand something that may not occur in models with higher dimensionality (there will be more ways to move around the trails left by Explorers). On the other hand, it is a phenomenon that it is hard to think of occurring in real scientific contexts.

So our findings for the medium range case are probably most relevant in applications. I said above that we can often observe a diverse set of both extractor-types and explorer-types in science. Our model could now be seen as showing the benefits of this status quo. But does it also provide an explanation of the status quo? We may think that the fact that the actual division of labour we observe is beneficial somehow is responsible for things being arranged in such a way that it occurs. But it is not obvious from our model why the most beneficial mixed groups of scientists should emerge. In our model, we simply picked different compositions of groups, and the agents could not ‘choose’ what type to be. But if we want an explanation of why beneficial division of labour may emerge, and if we think that scientists have some choice over how to behave, we should ask why an individual scientist should choose to behave more like an Explorer or more like an Extractor. We may think that Extractors are more likely to find *some* significant approach, while Explorers, if they find anything, will find the more significant approaches. In that case, more risk-averse agents would choose to be Extractors, and less risk-averse agents would choose to be Explorers. But our model reveals that this explanation does not work, at least not

straightforwardly. Recall again that for the medium ranges, not only is division of labour beneficial, but Explorers also do worse in terms of the average significance discovered at each move. So not only is being an Explorer more risky, it is also less lucrative in terms of average significance discovered than being an Extractor. So unless being an Explorer is intrinsically very attractive to some, it is not clear here why anybody should choose to be an Explorer, and thus how the beneficial division of labour should arise.

So far, we implicitly assumed that what society cares about (epistemic significance) is also what the scientific agents care about - after all, epistemic significance featured in their individual strategies. But if that is the case, being an Explorer does not seem to be an attractive strategy anymore, and what is in fact best for all cannot be achieved. So if society nevertheless wanted to achieve the beneficial diversity of approaches, what would be needed is some kind of 'bonus' for the kinds of discoveries Explorers tend to make. Interestingly, I think such a bonus system exists. Namely, relatively novel approaches tend to be rewarded more highly in terms of prestige or future research grants than less novel ones. My model suggests that these rewards for novelty could act as incentives for the less risk averse agents to behave like Explorers and explore approaches 'out in the unknown'. On that interpretation, novelty would not contribute directly to an approach's epistemic significance. It's not that novelty makes a finding somehow intrinsically more significant, or itself acts as a kind of theoretical virtue. Rather, it is a bonus that is rewarded to scientists, which makes sure that the scientific community finds all the most significant approaches - be they novel or not⁴.

These last reflections suggest that what I presented here not only provided a formal model of a kind of division of labour that already seemed plausible without such a formal model. We also learned something new and interesting: to make such division of labour possible, extra benefits need to be provided to explorer-type scientists, since they will make fewer, and, on average, less significant discoveries than others. In order to see whether and how different kinds of 'novelty bonus' could play this role, an interesting avenue for future research would be to build a model that includes such a bonus explicitly.

⁴ Thus, even though we have used a very different model, our results bear resemblance to Strevens' (2003) justification of the 'priority rule', namely the rule that only the first to discover a result gets all the credit. He used a marginal contribution model to show that the priority rule may encourage scientists to use their research efforts in a way that is socially optimal. More in line with the priority rule, rewards for the novelty of a result could be interpreted to be like rewards for being the first to have discovered something significant within a certain area of the landscape. This reward scheme, too, may help to ensure that research is conducted in a socially optimal way.

7. Conclusions

It is now widely acknowledged that science not only is a social endeavour, but that the social aspects of science may also play an important epistemic role. Kuhn (1977) pointed out the possibility that in science, there may be a conflict between individual and collective rationality: A group of scientists working in isolation, motivated only by the search for truth, may fail to produce as much scientific progress as they could. But if that is so, there is reason to think that the coordination of research efforts, and alternative (non-epistemic) incentives may in fact encourage scientific progress. The model I presented in this paper not only showed the division of labour between two types of scientists, both of whom are responsive to their social environment, to be beneficial. It also suggested that this division of labour will not occur unless there are some bonus rewards in place that introduce incentives for the individual scientists that do not correspond to the social value of their research. So we have found another instance of collective and individual rationality coming apart.

References:

- KITCHER, Philip, 1993. *The Advancement of Science*.
KUHN, Thomas, 1977. *The Essential Tension*.
STREVEENS, Michael, 2003. "The Role of the Priority Rule in Science", *Journal of Philosophy* 100 (2), pp. 55–79.
WEISBERG, Michael, and MULDOON, Ryan, 2009. "Epistemic Landscapes and the Division of Cognitive Labor", *Philosophy of Science* 76 (2), pp. 225–52.