Project Proposal: Explorers, Extractors, and the Benefits of Diversity in Science

Johanna Thoma
University of Toronto and Munich Centre for Mathematical Philosophy

1. Background and Research Question

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 point the way to fruitful areas of research, and the former 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 model seems to imply 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. The purpose of this research project is to develop a model that can provide such a demonstration. Using what I argue are more adequate descriptions of what I will call the 'explorer' and 'extractor' strategies, I develop a model that I hope shows the advantages of a mix of explorers and extractors over pure groups when it comes to exploring a field of research.

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. Given this insight, a literature has emerged that uses formal models to show that the incentive structures that in fact govern science, and the ways in which scientists in fact coordinate their work (the 'social structure of science') are indeed beneficial.

One such feature of science that has received special attention is diversity. Intuitively, diversity seems to be a good thing: A diverse set of research methodologies, questions, or prior beliefs in science may, for instance, protect us from biased results[1] or premature acceptance of theories[2], or furnish us with more robust results[3]. Furthermore, diversity may be a sign of *division of labour*, which economic theory tells us can lead to substantial increases in efficiency[4].

---

[1] There is a large literature in feminist philosophy of science on this point. See Anderson (2012) for an overview.

[2] See Zollman (2010), who explores the benefits of a diversity of dogmatic prior beliefs in early stages of research.

[3] Agreement of a variety of methods is often interpreted as making a result more reliable. See, e.g. Hacking (1983).

[4] Accordingly, some of the first models of the social structure of science (e.g. Kitcher (1990) and Strevens (2003)) were economic models of the cognitive 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 work in apparent isolation (which already amounts to a kind of division of labour), they fail to show that division of labour between explorer-type and extractor-type scientists is beneficial. I conjecture that this is because their 'follower' scientists end up duplicating the work of others much of the time, which has no added benefit in the setup of their model[5].

I will hence build a model where all scientists avoid merely duplicating the work of others. This makes independent sense for two reasons: First, research results, especially those of researchers working on approaches similar to one's own, are now mostly freely available, so we can assume that researchers already have access to the results from previously investigated approaches. 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. Still, we can meaningfully distinguish between extractor-type and explorer-type scientists even when all avoid duplication: Explorers 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. My central question is hence: Are mixed populations of extractors and explorers better at making scientific discoveries than pure populations of extractors and explorers respectively?

2. The Model

I model scientists as making discoveries on an epistemic landscape. Such a landscape represents different scientific 'approaches' with an associated 'epistemic significance' as patches on a grid. An approach is characterised by a methodology, a research question and set of background beliefs. Nearby patches in the landscape represent similar research approaches. I assume that approaches are very fine-grained: For instance, using the same methodology and background assumptions as somebody else in order to treat a slightly different research question counts as a different, but nearby approach. Epistemic significance is a numerical value which represents, roughly, the amount of scientifically or socially important results that can be obtained using a particular approach[6]. My model, like Weisberg and Muldoon's, assumes that epistemic significance is not distributed randomly on the landscape, but in two hills with single peaks. This makes sense when we assume that approaches that are similar to significant approaches are also likely to be similarly significant.

---

[5] Weisberg and Muldoon suggest that division of labour may end up being optimal 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, followers may not be worth their money if they produce no added value much of the time.

[6] See Kitcher (1993) on the concept of epistemic significance.

In epistemic landscape models, scientific progress occurs 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. I assume that all agents successfully determine the significance of an approach when they use it[7]. Scientists can move around the landscape using certain rules. The rules I will investigate can be described as follows in pseudo-code:

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: Move to the unvisited patch at the greatest distance to the neighbouring patches previously visited by other scientists. If several patches are at an equal distance, pick randomly between them. If there are no unvisited patches, 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 an extractor encountered other scientists before, she will have set her direction to get away from them, and she will go ahead unless she starts going 'downhill'. 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 unvisited patches in my Moore neighborhood?

> If yes: Go to the unvisited patch in my neighborhood that is closest to the previously investigated approach in my neighborhood with the highest significance (including my own previously investigated approaches). If there are several such approaches, pick randomly between them. In the first round, pick a random patch in the neighbourhood.
>
> If no: Go to the previously investigated patch with the highest significance. If there are several such approaches, pick randomly between them.

This rule expresses both a desire to make significant discoveries (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. This hence captures the mentality of scientists who, while concerned to make significant discoveries, want to stay close to what other scientists are doing.

---

[7] Though it may be controversial in some circumstances, I follow Weisberg and Muldoon in making this assumption.

I will compare how well populations of different composition but same size do at exploring the landscape when scientists are initially randomly placed. In particular, in a large number of runs each, I will compare the two pure populations with a number of different mixes of extractors and explorers. For that, I will use the following measure of success at exploring a scientific landscape: Let the total significance of the landscape be the sum of the epistemic significance of each of the individual approaches. What matters, and what I take to be the most important measure of success, is what proportion of that total significance the community of scientists discovers. I will compare the different populations with respect to how long, on average, it takes them to discover 10%, 20%, 30% and so on of the total significance and plot the results. I will also analyse and present graphically how discoveries unfold spatially in a few representative simulations, in order to understand how extractors and explorers divide labour.

3. Expected Results

I conjecture that a mix of extractors and explorers finds a greater proportion of the overall significance of a landscape faster than pure populations of explorers and extractors do, and hope to determine roughly what mix is ideal. I thus hope to show the benefits of a division of labour between explorer-type and extractor-type scientists when it comes to making discoveries within a given research field.

This result would be significant not only because it suggests that there is an epistemic benefit to the actual diversity of explorer-type and extractor-type scientists we find in science. I also take it to be plausible that the appropriate reward schemes exist to make both explorer-type behaviour and extractor-type behaviour attractive for different scientists. Research results tend to be more highly rewarded (in terms of reputation, or future grants), when they are relatively novel. The choice between explorer and extractor-type behaviour then resembles the following: Either researchers do work similar to what has already proven fruitful, in which case they are more likely to make significant discoveries, but these discoveries come with a smaller reward. Or they explore new territory, in which case it is more uncertain whether they will make significant discoveries, but if they do so, they will get a higher reward. Faced with this choice, it makes sense for less risk-averse scientists to follow an explorer-type strategy, and for more risk-averse scientists to follow an extractor-type strategy. If my model shows division of labour between extractors and explorers to be beneficial, this would then give an interesting spin to justifying the value we place on novelty in science: Rather than itself being an intrinsically significant aspect of a discovery, it may be thought of as a reward that makes possible a beneficial division of labour in the exploration of a research field.

References:

ANDERSON, Elizabeth, 2012. "Feminist Epistemology and Philosophy of Science", in ZALTA, Edward (ed.), *The Stanford Encyclopedia of Philosophy*.

HACKING, Ian, 1983. Representing and Intervening, Introductory Topics in the Philosophy of Natural Science.

KITCHER, Philip, 1990. “The Division of Cognitive Labor”, *Philosophy of Science* 87, pp. 5–22.

KITCHER, Philip, 1993. *The Advancement of Science*.

KUHN, Thomas, 1977. *The Essential Tension*.

STREVENS, 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.

ZOLLMAN, Kevin, 2010. “The Epistemic Benefit of Transient Diversity”, *Erkenntnis* 72, pp. 17–35.