

Social Epistemology and Individual Rationality

Draft: Please do not cite or circulate without Author's permission

Introduction

Prior to Kuhn, philosophers of science frequently saw scientists as archetypes of virtuous rationality. Scientists were viewed as machine-like seekers of truth. The data was simply there to be collected and interpreted, and each scientist worked to further knowledge in the world. But Kuhn shattered this view of the world, and spoke instead of communities of scientists, working within shared conceptual frameworks, working to resolve puzzles laid bare by this framework. This pursuit of normal science was in essence puzzle-solving, and Kuhn noted that scientists were motivated not just by the pursuit of knowledge in general, but instead they were motivated by a desire to have a reputation of being a clever puzzle-solver. Though a number of features of Kuhn's general account of science are mistaken, this was a fundamental insight for the philosophical literature. Scientists, instead of being theory machines, are instead self-interested agents. This is not to say that scientists have no interests in the pursuit of knowledge, of course. Scientists pursue knowledge, but *what knowledge they pursue* and how aggressively they pursue it are going to be determined not only by purely epistemic reasons, but also by social considerations, including incentive systems. The question then becomes how to leverage scientists' self-interest such that their cognitive labor results in maximum epistemic output. Both Philip Kitcher in *The Advancement of Science* and Michael Strevens in "The Role of the Priority Rule in Science" have offered models to analyze the role of incentives in knowledge production. While these models are compelling, we need to judge them on the basis for which they were developed: to better

understand the social role of scientific discovery, and what interventions that we might be able to make to encourage improvements in scientific output.

To do this, we must first consider what the considerations are in a program of social epistemology that differ from that of the traditional individual perspective. From there, we can consider what the goals were for the models Kitcher and Strevens present. In light of this, we will see the ways in which the models fail to adequately satisfy the desideratum. It is then that we can consider an alternative approach to the problem that is more suitable to the task at hand.

Social Epistemology and its Role in Science

If we examine the traditional questions posed by philosophers of science about scientific practice, we find a focus on questions like the relationship between theory and evidence, the model-world relationship, and a host of challenges within the context of justification of theories. Within this domain is a host of pressing questions and a great deal of invaluable research. But it is worthwhile to consider whether there are other questions that are equally pressing that have yet to receive the same amount of attention. One may notice that the common areas of research share several features. They all pre-suppose a particular problem is given to the scientist, outside of consideration. Frequently, there is an assumption of a single scientist (or lab) working on the problem, rather than a group of collaborators. Further, there is little analysis in the philosophical literature of the effects of the various institutions of science – conferences, associations, and journals surely have a role in information transmission amongst the scientific community. Finally, another notable absence from the traditional questions in the philosophy of science is the role that funding and funding agencies play in shaping science and the questions the scientific community pursues.

In one sense, it is surely reasonable that these considerations have been ignored.

They complicate any analysis of science tremendously. It is important to first get a solid grasp of the individual case of the lone scientist before we try and understand the workings of the scientific community. In addition, the tools to investigate the epistemological consequences of different social organizations is are different from, and almost surely less developed than, the tools that are more regularly employed by philosophers of science. However, both of these worries seem to now be overwhelmed by the promise of the gains to be made by investigating the social aspects of science, if only because we understand so little about it.

Social epistemology offers a potential research program that is huge in scope, and largely untouched by philosophers of science. While I will restrict the focus of this paper to a fairly minimal portion of it, I will draw attention to a few of the possibilities for research, much of which Kitcher mentioned in chapter 8 of *The Advancement of Science*. Staying with traditional questions, one can investigate how to incorporate evidence from other scientists into one's own beliefs. Related to this is how we might be able to determine an optimal rate of replication of others' results, and further, how we might be able to achieve this in practice. Kevin Zollman has done work on the trade-offs involved with how much scientists communicate with each other. Another direction in the same vein is an analysis of optimal levels of specialization in the scientific community, and how this might interact with communication structures. Finally, we can diverge further, and look to how we might be able to understand how scientific programs are chosen, and what room there is for intervention.

Of course, while social epistemological considerations offer several new degrees of freedom for philosophical investigation, it must include some constraints. Being sensitive to social concerns means that we incorporate non-logical considerations into our models of science, but we should not treat this as an attempt to undermine science's pride of place and to consider the contents of *Science* to be of no greater stature than the contents of *Maxim*.

The foundations must continue to be based on logic, but must also be expanded to incorporate rationality. We can continue to look to logical considerations to understand what the epistemic goals of science are, but we must turn to rationality to see what is possible to achieve.

Incorporating a theory of rationality into scientific considerations can be seen as sullyng the pursuit of science with mundane self-interest. It is not a claim, however, that scientists are only interested in money and fame. Rather, we should understand it as the recognition that most scientists are not purely selfless truth-seekers. They have real constraints, and trade-offs to consider beyond just what may be epistemologically required. The recognition of this does not necessarily represent a loss, however: it also points us to a means of encouraging or discouraging different lines of inquiry.

With this in mind, the focal question being considered for the remainder of the paper is that of project selection. If we imagine a situation in which there are two available projects, A and B, that both aim to provide a means to answer the same question, how ought a scientist choose which project to take on? Here we can imagine examples from history, such as how to go about determining the physical structure of DNA. One project might be X-ray crystallography, and another might be trying to build physical models. In *Advancement* Kitcher claims that the rational scientist working in isolation would choose whichever project has the greatest likelihood of success. If we extend this decision procedure to all scientists, we would find ourselves with a serious problem: all scientists would be working on a single project. There would be a great deal of duplicated effort, and our scientific bets would not be suitably hedged. If the “more likely” to succeed project in fact fails, then the scientific community would have to collectively start over.

This solution is clearly non-optimal, once we recognize that there are many scientists available whose labor we can divide amongst the projects. The question then becomes how

we can evaluate this kind of allocation in a way that is consistent both with our understanding of actual scientific practice and with what an optimal distribution might be. Kitcher presents a model that Strevens later develops, and it is Strevens' later model that I will consider. In Strevens' model, he seeks to demonstrate that the scientific community has already developed an optimal method for the distribution of cognitive labor, based on the priority rule.

The Priority Rule

Strevens seeks to offer an explanation of why the scientific community uses the Priority Rule as a reward scheme. Strevens argues that the rule has two components. First, scientists are rewarded based on achievement, not on effort. Second, only the first scientist to make a particular discovery is rewarded. (Strevens, p.3) Strevens considers the currency of these rewards to be prestige, not something monetary like salary or grants.¹ One curious feature of the Priority rule, which Merton calls a "dysfunction", is what Strevens calls the *Arago Effect*: even if two scientist independently arrive at the same results within days (hours, minutes) of each other, it is still the case that only the first gets any credit.(Merton, p. 322, Strevens, p.5) Rather than call this a dysfunction, Strevens argues that this effect is in fact a beneficial feature of the priority rule, for reasons we will see later.

Strevens considers three alternatives to the Priority Rule. First is simply an egalitarian distribution of prestige among scientists. Everyone who participates in the scientific community would then be equally prestigious. Second is rewarding scientists based on their effort. Hard working, ingenious scientists would be the most rewarded, even if their work does not produce many results. Finally, we have a model that is similar to the Priority Rule, except it would be a proportionate system of rewards – rather than the winner taking all, those who finish second would also receive rewards for coming close. (Strevens p.6) The goal of his

¹ Generally there is a strong correlation between prestige and salary, but this is not included in the analysis.

paper is to explain why the scientific community employs the priority rule instead of these other alternatives.

Strevens' Model

Strevens argues convincingly that the problem of optimally distributing scientists among projects is equivalent to a resource allocation problem. That is, there is a certain good (cognitive labor) that can be allocated amongst different consumers (scientific projects). These projects have marginal utility curves (success functions) representing their ability to productively utilize additional resource. These success functions output probabilities of success. Each project has an associated value if it is successful. This sort of model is common in economics and is known as a constrained maximization problem. The work to be done in the model is to determine what method of allocation produces the best epistemic results – which reward system creates the most successful scientific projects.

Strevens assumes a number of things in all of his models. First, he assumes that scientists are rational, which here can just be taken to mean that they are utility-maximizers. Simply put, scientists respond appropriately to incentives. Second, he assumes that there are two projects at a time for scientists to choose from. Third, he assumes that each project has an intrinsic success function, which takes as input units of cognitive labor (time from scientists), and outputs objective probabilities of success. The features and values of the success functions are common knowledge among the scientists in the model. (Strevens, p.10-13)

Strevens proposes three different reward systems that serve to function as allocation mechanisms in the modes: *Merge*, *Goal*, and *Priority*. These are designed to be tests of the

Priority Rule and its alternatives.² Under *Marge*, each scientist is rewarded based on how much she contributed to the increase of a project's probability of success. That is, the reward function for a given project is defined by the derivative of its success function. The rewards are strictly decreasing, since the success functions are concave.³ Receipt of the rewards are not distributed based on actual success of the projects, but rather based on the contributions of scientists. As Strevens notes, scientists under this model scientists are rewarded based on both their effort and crucially, their timing. Those that start projects contribute more to the marginal probability of success, and so are rewarded more. (Strevens, p. 18) This encourages scientists to “get in on the ground floor” of a project. This method of allocating and rewarding resources results in an optimal distribution of cognitive labor, provided that the projects that scientists can choose from are independent – the success of one has no impact on the success of another. Unfortunately, this is frequently not the case in science, and *Marge* will produce sub-optimal results in those instances. (Strevens, p.19-20)

For example, in a “winner takes all” case, in which there are two projects that are competing to make the same discovery, Strevens points out that the optimal distribution would allocate even more resources to the project with a higher intrinsic value than what *Marge* does. So, *Goal* is introduced as a modification of *Marge*. While it maintains the character of *Marge*, it varies in that it does not reward projects that are unsuccessful. It only rewards scientists who contribute to projects that succeed. To achieve the distributional shift of cognitive resource to projects with higher intrinsic probabilities of success, Strevens appeals to risk aversion. Someone who is risk averse is willing to pay for a reduction in risk. So, as Strevens puts it, they have a decreasing marginal utility for a given good (in this case, prestige). A standard example of this phenomenon is that an agent would prefer \$100 to a

2 Strevens does not consider an egalitarian reward scheme, presumably because it would do nothing to solve an allocation problem – scientists would be indifferent between projects.

3 The technical details of the models will be elaborated on in the criticisms below.

10% chance of winning \$1100, even though the expected benefit of the second lottery is higher. So, if agents have suitable levels of risk aversion, *Goal* will cause the distribution of scientists to shift in a manner that accords well with the optimal allocation. (Strevens, p.25-27)

Strevens argues that if we switch from *Goal* to *Priority*, in which scientists are only rewarded if their project succeeds *first*, then we can increase the distributional shift that *Goal* began. This is because the only shift in reward distribution occurs when both projects succeed, as only the first project to succeed is rewarded. Because of this shift, signing on to either project offers a lower expected utility, but Strevens argues that it is most likely that the project with the higher intrinsic probability of success will also be the faster project. As such, the expected utility of joining that project will decrease by less than the expected utility of joining the other project. So, more scientists will work on that project than they would have under *Goal*. Because of this, the Priority Rule has the property that it will maximally benefit society by suitably biasing resource allocation to those projects that are most promising, which still hedging epistemic bets by providing resources to lesser lights. (Strevens, p.29-30)

A Note About Criticisms of Simplifying Model Assumptions

In the following sections, I will consider a series of objections to the assumptions in Strevens' model that ultimately result in my rejection of the model. Before I begin, I would like to note that each of these assumptions have precedent – in particular, many of these are standard assumptions made in economic models. It is only relatively recently that economists and others have begun investigating models that weaken these assumptions, most notably in the context of agent-based modeling. Indeed, most economic models continue to be population-based models rather than agent-based models because of a sharp increase in complexity for the newer modeling strategy. Analytic proofs become fairly challenging with

the additional complexity. It is also extremely common in economic models to examine binary choices (or only two agents), as going to N choices can get analytically intractable. As such, Strevens is not doing anything out of the ordinary for this sort of modeling.

While simplicity in models is most definitely a virtue, models can be too simple. In particular, since Strevens uses the model to defend his argument that the Priority Rule has certain optimality properties, such that it would have emerged to beat out other incentive systems, it should be capable of defending the position, ideally by taking the history of science into account. The actual practice of science seems to deviate from the assumptions of the model in important ways, such that his modeling results are not clearly about the scientific community. Further, I offer reasons for why these deviations are such that if he changed his model to accommodate them, his results would be qualitatively different. I also provide evidence that this is the case by means of simulations that weaken these assumptions. Finally, I suggest that the style of model is insufficiently explanatory for the questions he asks. It is in these sorts of cases that models can prove to be too simple.

Some Difficulties with the Model

There are five basic problems with the Kitcher-derived model that Strevens employs. The first three are epistemological assumptions. First, Strevens' model includes the simplifying assumption that all agents have common knowledge of the objective "intrinsic" probabilities of success for any given project. Second, a consequence of the first assumption is that we have some basis for being able to accurately evaluate the probabilities of success for a project, even before we have begun any work on it. Third, the model assumes that all agents have common knowledge of the current distribution of cognitive labor. The fourth assumption is that there exists an appropriate ratio of available scientists and available work. Finally, the model to support the priority rule can be seen as requiring agents to have

inconsistent attitudes toward risk. Let us briefly look at reasons for why these assumptions might pose a problem. Then we will turn to simulation evidence that support these reasons before turning to an alternative approach.

The First Two Assumptions

The assumption of common knowledge of objective probabilities faces a number of difficulties. To begin with the most fundamental assumption, we see that each project does have as an objective property a probability of success. That is, each project has a specific probability of achieving its goal. This is an intuitive assumption – after all, it is perfectly reasonable to claim that using properly calibrated telescopes to observe the motion of stars will be more successful at determining the age of the universe than would consulting the Old Testament. Some projects and methods are simply better than others, regardless of what we might believe. There may be some population that has continued to work from an assumption that the Earth is flat, but that certainly does not make it so. Those of us who work from assumptions that include a spheroid Earth are going to be more successful. Of course, these are extreme cases, but it is not a large jump to say that more subtle cases are a common occurrence. But we can already begin to see here why the assumption is problematic in the end. Let's assume that each project has an associated probability. Further, let's assume that it has a probability function that takes as input the number of hours spent working on the project, and outputs the probability of success (which is what is needed for the model). But then, we find some difficulties in the history of science. . That is to say, phlogiston theory had a non-zero probability of being a successful theory of chemistry. Alchemists had a non-zero probability of turning lead into gold. Phlogiston theorists may have believed that they were right, but the cards were stacked against them, no matter how many scientists they added to their ranks. It must be the case that false theories have a zero chance of success, at least as

a metaphysical property. It may be the case that they can serve as partial explanations, or have some other epistemic virtue, but metaphysically they must have a zero probability of success. But we find that these facts were not common knowledge. Lots of people devoted a great deal of time and effort to working on projects that we now know to be fruitless. Most of science has this character. So, even if we grant that any given scientific project has an objective “success function” we seem to be on rather weak ground in assuming that individuals have any sort of epistemic access to this, let alone much agreement amongst each other.

The Third Assumption

The model assumes that agents know the current distribution of labor amongst all scientists. Even if we assume that the model represents all and only those agents who are engaging in a niche area of research, it still seems difficult to claim that they know precisely how many agents are working on each project. This is necessary to be able to accurately calculate what the marginal returns are for working on a particular project. It is more reasonable to suspect that scientists have knowledge of who is working on which project amongst those scientists with which they regularly communicate. But this is a weakening of the initial assumption.

The Fourth Assumption

The model assumes that there exists an appropriate ratio of scientists to projects that require resources. That is, there are not too many scientists to be useful, and there are not too few scientists to make any progress on a given project. This can be investigated in a case in which there are more than just 2 projects, but it can also arise from projects that are particularly difficult and require a great deal of labor. If there are too few scientists, they will

distribute themselves in a way that produces no result.

The Fifth Assumption

The model for Priority that Strevens employs depends on a notion of risk aversion that results in more agents working on the projects that are more likely to succeed. However, the distributions of awards that the rule itself generates are much more risky than alternative rules. So if agents were in fact as risk averse as Strevens suggests, they would be much less inclined to want to choose the distributional rule in the first place. As such, it seems to be at least a contentious claim that the agents are risk averse as he describes them to be.

Simulation Setup and Results

Along with Michael Weisberg, I created a series of simulations to evaluate the weakening of these assumptions. In each, I implemented one of the main distributional rules that Strevens discusses – Marge, Goal and Priority. These were implemented as agent-based models in NetLogo. Strevens never specified what the precise shape of the success functions should be, but following Kitcher we treated success functions as logistic growth equations. In the models, agents are spatially distributed randomly across a torus, and assigned a random initial project. Each agent has a radius of vision, which allows the agent to see all the agents within this radius, and which project they are working on. Each round, agents use the information available to them and determine which project would be most advantageous to work on. It is important to note that these simulations maintain as many assumptions as possible from the Strevens model. Just as in Strevens, there are no dynamics involved. As such, the simulations reach a steady state within a few rounds.

[Insert image of simulation environment]

If we look at the simulation results of the weakening of the model assumptions, we see

that the distribution is skewed with each weakening. Most prominently, however, is the first assumption. Once we weaken this assumption by shrinking the radius, we find a number of clearly troubling effects. Notably, some projects simply die out – no agents work on them. In other conditions, the distribution is the reverse of the optimum. This is clearly a crucial assumption.

[insert several graphs and brief explanations here]

In light of this evidence, it is necessary to decide whether this kind of model is worth pursuing more. As we incorporate changes into the model to correct the problematic assumptions, the model becomes increasingly unwieldy. The models incorporate several dozen parameters, just to address the initial question. As a platform for further research, it does not seem as if this is the most fruitful avenue.

A New Model

A challenge in developing a more useful model of scientific research is that we need to be able to incorporate an idea of individual decision-making and a social context for those decisions while recognizing that there are facts of the matter to discover. While social forces may influence what individuals work on, they do not influence what is true. Incorporating these concepts into a single model is challenging – Kitcher and Strevens make significant efforts in their models. We have some aspects of a decision-theoretic problem, and some aspects of a game-theoretic model. But at the heart of it, Kitcher and Strevens were right to think of this area of investigation as an optimization problem.

Rather than frame the model as a constrained maximization problem, we can instead think of it as a variant of a hill-climbing problem. In a hill climbing problem, an agent or agents who have no knowledge of the landscape try and find higher ground. The simplest way of

doing this is by means of random moves. Each agent moves randomly, and then sees if the elevation is higher than before. If so, they continue with random moves. If not, they back up and try again. As we allow time to go to the limit, all agents will arrive at a local optimum.

This model immediately grants us several features. First, hill-climbing problems are naturally represented in an agent-based context. There is a separation of what a given properties of an agent are (including its beliefs about the world) from the properties of the objective world. This allows us to maintain each scientist as an individual decision maker, and allows us to introduce heterogeneity without significant difficulty. That is, different agents can adopt different hill-climbing strategies, whether the differences be based on their method of inference, or their utility function. Second, imagining the model as a hill-climbing problem also grants us a clean separation of the distinct features involved in the model. We can investigate the landscape independently of the agents upon the landscape, so as to separate out what is the case in the world from what agents think about it. In this way, we are able to better position ourselves to look at the social process of scientific inquiry.

There is more that is needed than just changing the model, however. We need to also re-calibrate the questions we ask in order to more accurately get at the social aspects of science. To do this, we must generalize the question and Kitcher and Strevens raise. Part of the problem with their analysis is that they in some ways mistake where one should focus on the individual and where one should pay attention to the social group. Their model is from the perspective of the “representative individual” that is commonly employed in economics, and this individual uses information about the aggregate decisions of the group. This washes away individual differences, and also does not allow for the investigation of social dynamics.

Rather than frame the problem of project selection as a single choice between two projects, and look at how scientists that are aware of other scientists ought to make their decisions, we can instead consider longer-run dynamics of scientific research. That is, we

can consider the scenario in which scientists make a series of choices through time, and how these independent choices can affect others. In this way, we can represent diverse agents whose decisions are influenced by both the landscape and their peers.

To do this, we interpret the landscape as being a conceptual landscape, in which each point on the landscape represents a possible research program. To unpack this concept, think of each point in the landscape as being the unit of research that is represented by a potential paper. Points that are spatially close to one another share common features, but do not fully overlap. In this way we try and interpret the spatial relationships of the landscape as indicative of conceptual relationships between research programs. In this way, the spaces near where an agent currently is are suitably related to the current project such that it is reasonable to expect that a scientist would move to one of them at some point in time. From this, we can think of a path across this landscape as being a long-term research agenda.

While we have discussed how the directional features of the model are interpreted, we have not yet discussed what elevation means. As it is a hill-climbing model, we are left with the implication that higher elevation is better. To maintain this as a meaningful notion, we interpret the elevation of the landscape as the significance of the potential result, where we understand “significance” only in terms of epistemic value. Following Kitcher, we assume that the challenge isn’t discovering truth, but instead discovering significant truth. “Significant” truth is admittedly not a fully precise notion. Epistemically significant results are going to shift over time to some degree. Calculus was a great discovery, but now high school students learn it, and college students learn how to prove portions of it. So in some ways one might say that calculus is less significant now, if we think of significant as being important and suitably impactful. We can largely avoid this problem for now by simply interpreting the timescales to be involved in the model as being on the order of years rather than centuries. This notion can be extended, or potentially augmented, as will be discussed briefly later.

So, given our interpretation of the landscape and what it represents, the base model then interprets the agents as employing a hill-climbing strategy to seek out the most productive research agenda they can find. As mentioned before, by separating out the landscape from the agents who inhabit it, we make it easier to investigate the potential effects of social interaction and heterogeneity on behalf of the agents. It is also conceptually more clear to see that individually-rational agents can understand themselves to have both common and individual goals. Simple questions can arise from the base model, such as whether we want to optimize our reward strategies for coverage of the landscape, coverage of parts of the landscape, or optimal time to getting someone to the highest elevation, or even optimizing the time it takes for some portion of the population to get to the highest elevation. Different landscapes might encourage different strategies, and can address different kinds of questions. For example, Strevens' project could be re-interpreted as being a landscape with two hills, and we can look at what agent behaviors are best to encourage maximal coverage of significant space, versus optimizing the time it takes to reach the global maximum.

Some virtues and extensions to this approach

The model of science as a landscape to be traversed is one that has a number of features that aid in philosophical investigation. As has been mentioned, the model separates what is properly strategic activity – the cooperation and competition between individual agents seeking to maximize their utility – from what are properly conceived of as features of the world – the particular facts of the matter that are being investigated. This is a significant modeling convenience, and allows independent changes to be made in a much more straightforward way. It also helps to ensure that each intervention in the model is properly targeted. Extensions to the base model do

not require invasive changes to the codebase, and most importantly, the interpretation of the model remains much more consistent.

By adopting this alternative modeling approach, we are also able to avoid a basic problem that is difficult to resolve in the Kitcher-Strevens approach. Namely, though Kitcher explicitly rejects the idea that there is or could productively be a central planner for science in *Advancement*, the model that is eventually used is no different from how a central planner would do it. The information conditions are such that each agent knows everything there is to know. As Kitcher was right to want to reject a central planning model, we want to ensure that we are able to avoid building in any such assumptions into any model of science. By adopting an agent-based model, it is straightforward to ensure that this is not the case. Each agent has only limited information about its surroundings, and the information may be different between agents. No agent knows everything, nor can agents directly influence the behavior of other agents. Because of this design, it would be extremely difficult to accidentally fall back into a central planning model.

In virtue of some of these other features, the model is extremely general. It does not rely on many strong assumptions. The strongest is that there is a fact of the matter about scientific outcomes, and some outcomes are more significant than others. Because of this generality – after all, it is just an optimization model – there is a great deal of flexibility in the model such that extensions are fairly straightforward. If it is ever the case that this becomes a more active area of research in epistemology and the philosophy of science, working from a common platform helps to ensure that the modeling assumptions are well understood. It will eventually be in a position to help lower the bar to entry for research into this area, as new extensions do not have to start from scratch. In this way, smaller questions can be more readily investigated, as they require much less effort.

Though research using this model is very much in the early stages, several extensions to the base model have begun. I will briefly mention two of the most interesting, to provide examples of the utility of the model itself. As discussed before, a potential worry with the interpretation of the model is what exactly “significant” means. There needs to be some interpretation of the landscape, and so we want to ensure that it is sensible. But, there is no reason why we have only one landscape, or even that each agent necessarily sees the same landscape. We can imagine that the model has several overlapping landscapes, each of which represents a gradient of some possible desideratum. For example, one might represent funding opportunities, while another represents application to a given industry. Agents then can decide whether or not they care about each of these landscapes, and construct their own internal landscape based on some linear combination of the “actual” landscapes. It is here that we can see how some measure of cooperation between agents can be extremely useful. We can also then interpret different research agendas as partially determined by a particular view of the landscape.⁴ Some scientists might be more motivated by money than others, maybe some by religion, and perhaps some just for the love of wisdom. This multiple landscape model can help understand how these different motivations can be productive in the aggregate.

Another extension is to return again to the idea of significance, and have some means to resolve the problem that significance may change over time. Results surely become less important once someone else has already found them. To represent this

⁴ This idea is in part inspired by work done by Scott Page across several publications, including *The Difference*, Princeton University Press 2007

idea, we can borrow ideas from ecological modeling, and make the agents be able to change the landscape itself. A straightforward way of conceiving of this is that we can imagine the landscape as being a mound of ice, and each agent is a heat source. As the agents traverse the landscape, and perhaps linger on given locations, they melt the ice, changing the landscape itself. This representation allows us to look at longer time scales, and begins to provide a mechanism for thinking about how new questions might be shaped by the research that we engage in. What becomes significant is to some degree path-dependent. Some questions become important due to the fact that they are needed to resolve questions raised by previous research. With the “melting” model we have a rudimentary way of investigating the dynamics of the changes to the intellectual landscape.

A Return to the Question: Why Social Epistemology?

As we have seen, social epistemology raises a number of new technical challenges that do not exist otherwise. We have to learn to deal with interactive rationality, we have to maintain several desiderata at once, and we potentially lose some of the analytical exactness of proofs that we might be more comfortable with. But we have much to gain by taking on these additional burdens. New questions open up to us, and we also gain new insights on old questions by expanding our field of view.

Most importantly, however, is that it is just the case that insofar as we are epistemic agents, we also social agents. We collaborate on research, we accept testimony from others, we go to conferences and read journals. We not only try and

learn from the investigations of others, but we take part in institutional structures that help to ensure that the common goals of epistemic discovery are able to be met. So occasionally we find ourselves refereeing papers, commenting on them, and reproducing results. Changes to these procedures and institutions will result in changes to our individual practices. So it seems that as we put effort into better understanding what individuals ought to do to improve the quality of their beliefs, we should also put forth effort to understand the social dynamics that can collectively improve our outcomes. The models by Kitcher and Strevens were exciting beginnings to this line of research, and pointed to new directions in epistemology and the philosophy of science. The model I have advocated will hopefully be able to further promote the investigation of these questions in a more fruitful way.