Increasing returns in science

A model of the dynamics of scientific activity

Rogier De Langhe & Matthias Greiff

Rogier De Langhe
Centre for Logic and Philosophy of Science
Ghent University, Belgium
Rogier.DeLanghe@UGent.be

Matthias Greiff
Department of Economics
University of Bremen, Germany
Tel: ++49 421 218 3595
greiff@uni-bremen.de

The authors are grateful for helpful comments from the participant of the TILPS Workshop "Formal Modeling in Social Epistemology", October 2008, Tilburg University (The Netherlands), as well as comments from Wolfram Elsner and Torsten Heinrich.

Increasing returns in science

A model of the dynamics of scientific activity

Rogier De Langhe & Matthias Greiff

1. Introduction

Social epistemology is characterized by a recognition of the social dimension of knowledge acquisition. Allegedly, there used to be a time of individual geniuses reading the book of nature. Alas, no more. Contemporary science has accumulated to such an extent that no scientist, however genial or productive, can collect all evidence on his own. Scientists have started to specialise, depending increasingly for their evidence on each other rather than the world. In this *distribution of labour* in science, scientists can do no other than to *trust* each other's *testimony*. Science has become a network and the scientists merely its nodes. As such, our main focus is the network and not the nodes. A key question that arises once a cognitive effort becomes social is the extent to which this specialisation should be pursued. Helen Longino characterised this question of the division of labour as "*the question whether and when to pursue research that calls a community consensus into question or to pursue research that extends the models and theories upon which a community agrees.*" Communities of epistemic agents need to find a certain balance between specialisation and diversity. An important focus for our model is to detect the mechanism by which this balance is generated. Knowledge of and control over such a mechanism could provide us with a key instrument for institutional design.

The model we propose does not intend to prescribe what rational science *should* look like, but tries to grasp the dynamics of actual scientific practice. We believe that understanding the present state of affairs is likely to be conducive to attaining the state of affairs which is prescribed by more ideal models. As such, our model should fit with episodes in the history of science. The epitome of such a historicist approach is Thomas Kuhn's 'Structure of scientific revolutions'. This approach does not focus on how individual beliefs are formed and supported, but how these beliefs add up to entire schools of thought, 'paradigms' or 'research programs' and their interaction. This works through in social epistemology. An assumption from classical epistemology which has come under fire in social epistemology is that of epistemic individualism. Knowledge acquisition is a distributed effort in which agents need to rely on testimony from others. "By accepting each others' testimony, individual researchers are united into a team that may have what no individual member of the team has: sufficient evidence to justify their mutual conclusion." (Hardwig 1991, 6) This brings the institutional environment (which constitutes the interactions of epistemic agents) to bear on the conclusions reached. Hence, also this institutional environment and its history becomes part of the domain of study by epistemologists. Hence, a challenge for our social epistemology model of scientific activity is to meet Philip Mirowski's concern that *"A relevant congenital tic of the American philosophy profession (although, it must be conceded, not its alone) is a demonstrated unwillingness to regard science as an historically changing entity, not just in the realm of epistemic ''values'' but also in terms of actual social structures."* (Mirowski 2004 , 285, our emphasis)

The importance attributed to social and historical factors, the theory-ladenness of data and the idea of incommensurability tends to estrange it from those scholars interested in devising ideal accounts of science. It leads to a number of puzzles: Why do scientists sometimes not update their belief upon receiving evidence that conflicts with their beliefs? Why does research tend to cluster? If dissent is irrational, why is disagreement a persistent feature of science? On the one hand, dissensus seems to be omnipresent. Whether it is quantum mechanics, international relations theory or indeed even forest management, diversity and dissensus is ubiquitous across the spectrum of the sciences. Even highly formalized sciences such as logic and mathematics are divided into different schools of thought, debating fundamental issues such as the acceptability of certain kinds of inconsistencies or the existence of numbers. Others argue that there is an ever growing body of scientific results on which a consensus formed; it seems only a matter of time until all dissent has disappeared. “The positive argument for [convergent] realism is that it is the only philosophy that doesn't make the success of science a miracle” (Putnam 1975, 73). The model should explain this succes, both of results and of number of adopters. But Larry Laudan (1981) compiled a long list of once successful theories which are now ridiculed by the scientific community; and why would it be different this time around? In other words, we want a model which can both explain the success of science and pessimistic meta-induction. Our model has to account for three things: the existence of dissensus, the emergence of consensus and the dissolution of that consensus. This is a tough challenge, as Larry Laudan himself noted: “[S]tudents of the development of science, whether sociologists or philosophers, have alternately been preoccupied with explaining consensus in science or with highlighting disagreement and divergence. […] neither approach has shown itself to have the explanatory resources to deal with both.” (Laudan 1984, 2)

One of the most influential accounts concerning the distribution of labour is science is ‘The division of cognitive labour’ by Philip Kitcher(1990). Our model aims to extend this line of thought. The basic problem of that paper is the so-called “CO-IR-discrepancy”: the mismatch between a scientist’s individual rationality (IR) and the ideal balance between specialisation and diversity, the community optimum (CO). If scientists were all to pursue the same path, namely that which is best supported by the available evidence, then there is no diversity and the community optimum is unlikely to be reached, provided that, as Kitcher assumes, full specialisation is undesirable. Kitcher solves the discrepancy by introducing social and other factors, such as greed and stubbornness, which scatter scholarly attention and thus bring diversity into the scientific community. He then reformulates rationality from pursuing the problem-solving method which intrinsically has the best prospects of success irrespective of what others in the community are doing to choosing to belong to a community in which the chances of being the first to discover the correct answer are maximized. The latter case takes into account the distribution of research effort already present in the research community: prospective individual returns *decrease* as the number of scientists following a certain path rises. As a consequence, it becomes rational for the individual to pursue diversity, thus solving the CO-IR discrepancy.

We believe that Kitcher has made a valuable contribution in framing the problem and presenting a solution, but claim that his results have only limited scope. Kitcher only considers the case of choice between two research methods in search of a definite, true answer. The result is the modelling of a scientific community as a closed system, with a definite ending point and decreasing marginal returns as the endpoint nears (Kitcher 1990, p.12, footnote 8). Kitcher calls this an innocent simplifying assumption; understandably, because this way of simplifying is standard practice in neoclassical economics. We think, for a number of reasons listed below, that the model will not be robust in the face of concretization of its idealizations. This robustness is nevertheless of utmost importance, because if the scope of the model cannot be broadened to actual scientific research, where well-defined puzzles and ready-made evidence are often not in store, then Kitcher's model is too weak to extend its conclusions to the division of cognitive labour in scientific communities.

1) The development of a theory or research programme is an open-ended process. Its boundaries are not set and there is not a point at which it can be said to be fully developed and confirmed. Kitcher simplifies the problem as a choice between two problem solving methods of which one and only one will bring the solution. This is not an idealisation but an entirely different ball-game.

2) The literature learns that the applicability of terms like 'truth' and 'solution' is problematic in the context of research programmes. Yet it is these very concepts which smuggle in the decreasing returns in Kitcher's model because they cause decreasing marginal returns as the endpoint nears. Leaving these concepts out means we can't take decreasing returns for granted anymore.

3) Kitcher's system sees a research community as a set of individuals, each individually possessing the ability to come up with the entire solution of the problem. In reality however, as described at length in the social epistemology literature, knowledge is distributed and no individual is capable of developing an entire research program from scratch (everybody is, with Newton, "standing on the shoulders of giants"). Individual scientists have to trust each other, hence the problematique of *testimony*. This leads us to model scientific communities as networks, where it's the network and not the nodes that count. A crucial feature of networks is that they exhibit increasing returns: scientific communities, similar to networks like the internet, increase exponentially in strength every time a node is added to the network. As such, the addition of one scientist to a scientific community doesn't decrease the other scientists' prospects of success (as in Kitcher) but increase the prospect of success for the community exponentially. In brief, Kitcher's model does not explain why scientists try to persuade each other. With decreasing returns to adoption, it is not rational for a scientist to get other people to join the path he is pursuing.

4) Information is characterized by falling marginal costs (increasing returns to scale means falling marginal costs) because once a unit of information is produced (an idea, a book, a score,...) it can be distributed at virtually no cost. As such, the more people use it the cheaper it becomes and marginal costs keep on falling indefinitely.

We propose a model which has as its starting point the proposition that scientific knowledge is characterized by increasing returns, while Kitcher's model assumes decreasing returns to adoption. The difference between assuming increasing instead of decreasing returns is not trivial; indeed it has profound effects on the resulting system. As will be shown, this makes it an entirely different ball-game. The purpose of our version is to make a model that is more robust in the face of de-idealization than Kitcher's and as such can be extended to bear on the problem of division of labour in science.

2. Increasing returns to adoption

In response to the shortcoming of Kitcher's proposal identified in the previous section, we want to suggest a dynamic model characterised by increasing returns to adoption. For this purpose we draw on a series of papers by Brian Arthur, initially designed for problems of technology adoption in network industries, in which he proposes a formal way of capturing the dynamics of systems exhibiting such increasing returns to adoption. In his opinion, many increasing-return problems fit a general nonlinear probability schema. Using Arthur's work as a starting point, the following analogy between science and network industries is then proposed. Let's say a scientific discipline is like a table to which scientists add contributions, and the probability that a given scientist will add a contribution (say, a paper) to a certain pile (or 'network' or 'cluster') on the table depends on the share this pile of contributions in the total number of contributions on all the piles laying on the table. This introduces increasing returns to adoption, because the more people contributing to a pile, the bigger the odds that a new contributor will contribute to that pile.

This is a specific response to what Goldman (2001) has coined the 'novice/2-experts problem'[1]. Although responses to this problem differ, we follow Hardwig (1985, 1991), who holds that laymen are basically *blind* when having to choose between experts. Precisely because he is a novice, the novice has not yet spent any energy in getting to know the field, so he can not rely on his knowledge of the different networks, he cannot assess the reliability of different experts and he has no oversight of the discipline which would enable him to compare the available evidence for different networks or to know which network has the most adopters. Instead, what we propose is simply that the probability of the novice to contribute to network x is equal to the share of x in the discipline. This can be interpreted as that the novice is likely to adopt to those networks to which his peers adopt. And the odds that those peers will be members of network x is equal to the share of x in the discipline. So what is taken out of the equation here are the irrealistic assumptions which require agents to have perfect oversight over the discipline or some prior knowledge about its contents (which would make novices not so novice after all). In a novice-expert situation, the novice is 'blind' because no prior investments in extensive training and special competence have yet been made, the novice faces huge costs in order to assess diverging claims. In this situation, the most rational way to proceed is to take the way of least costs, i.e. drawing on the closest available experts (such as teachers). At this point, it *does not matter* whether the choice is right or wrong (i.e. the expert's views are true or false), because for the novice the expense is simply too high. This is an essential feature of our approach: the most important is not doing it right but doing the same. The rational tradeoff to be made here is to prefer compatibility because of the premium networks attribute to it ('network externalities'). Engaging in a certain amount of specialisation (i.e. make a choice, even if one is not fully convinced; pursuit without belief) will quickly outperform the option of suspending judgment and getting to know all available angles and weigh their respective body of evidence, including acquiring the information necessary to weigh this information accurately. Turning to that cluster which happens to be the cheapest one available at the time of entry, in terms of previously acquired skills, topics of interest, etc. is the rational thing to do, because it allows them to spend time using the methods they master best, answering the questions they themselves deem most relevant,... The point is that it doesn't really matter which cluster is chosen, as long as a choice is made. A choice (and the resulting specialisation) is needed in order to develop a certain programme thoroughly, devise the best arguments for it, come up with critical tests and eventually perhaps even allowing others to falsify it in the future.

Once this initial choice is made, the next choice on which cluster to contribute to no longer simply depends on the share of the cluster in the discipline. Every contribution counts as an investments in acquiring specific training and special competence. Through the lens of her contributions, the scientist will learn more about the available evidence and arguments. In principle, once a scholar is in the field he can change the cluster he contributes to with every contribution he makes. This might lead the novice to assess that his initial choice of cluster

---

[1] 'The novice/2-experts problem is whether a layperson can justifiably choose one putative expert as more credible or trustworthy than the other with respect to the question at hand, and what might be the epistemic basis for such a choice?' (Goldman 2001, p. 92)

was not such a good one after all. However, as investment progresses, sunk costs rise and the rational option for a scientist at this point is not simply to switch approach every time he regrets his previous choice(s), e.g. because of being confronted with evidence conflicting with his views, but to weigh this against sunk costs. As such the choice which cluster to contribute to is path-dependent. Furthermore, although not exclusively anymore, the relative share of the cluster still matters. Larger clusters mean more contributions have been made, resulting in better quality of the arguments and evidence produced in it. As such the positive feedback mechanism keeps on operating once a scholar is in the field.

In both episodes, as a total laymen and in growing to be an expert himself, *transaction costs* such as limited information, limited cognitive abilities and limited oversight weigh heavily on what is rational for the scientist to do. These costs have to be taken into account. They are not simply ‘noise’ but constitute an essential reason for engaging in science in a networked fashion in the first place. They cannot be idealized away without also dispensing with the distributed nature of scientific practice. Hence, ignoring them would render our results unfit for social epistemology. However, and tragically, this implies that *rationality* will not necessarily be truth-tropic, not on the individual level nor at the aggregate level. Although there is a lot of room for evidence in our model, rational agents in a network will not only be driven by evidence, but also by network externalities. As such we believe that, for a formal model in social epistemology, network externalities constitute an essential part of scientific practice. Lone experts, without any form of epistemic dependence, will simply not be able to come even close to the level of sophistication made possible by dividing labour among scientists. The social dimension of knowledge is not only restrictive but also constitutive.

Networks are built around a certain *standard*. This can be a shared language, protocol, ... To apply this to scientific communities, a standard can be taken to mean a shared core of basic concepts and theoretical assumptions, resembling what can be seen as the ‘core’ of a research programme. Scientific contributions form *clusters* around these standards. The study of a scientist working within a certain research programme and the research programme itself could then be approached in the same way as a computer user adopting to an operating system (e.g. Apple, Windows,...). This is the analogy used to apply Arthur’s formalism to the problem of the distribution of labour in science. In what follows we elaborate on his formal characterization of technology adoption under increasing returns in order to extend his analysis to the dynamics of scientific practice. A variant of this approach can be found in Zamorra-Bonilla (1999).

3. The model

Let us consider a population of N epistemic agents. There are J different clusters. Let $\overline{P}$ be a $J{\times}N$ matrix, its n-th column $\overline{p}_n$ being agent n’s (a priori or intrinsic) preferences for all clusters. There is a $J{\times}1$ vector E, its j-th element being the available evidence for cluster j. We normalize E and denote the vector of relative available evidence by $\hat{E} = (\hat{E}_1, \hat{E}_2, ..., \hat{E}_J)'$. Let there be a vector C being of length N, its n-th element $c_n>0$ being the weight agent n assigns to relative available evidence. For the moment assume that $c_n=c$, that is c_n is the same for all agents. As mentioned above, knowledge acquisition is a distributed effort where agents have to rely on each others’ testimony. A large value of c means that an agent has high trust in her fellows’ work and places much weight on compatibility. We interpret c as a proxy for trust, telling us how much weight each agent puts on evidence produced by colleagues. Let the *likelihoods of pursuit* for each agent n be given by equation (1).

$$\pi_n = p_n + c\hat{E} \qquad (1)$$

We define the following decision rule: Make a contribution to the cluster with the highest likelihood of pursuit (i.e. the largest element of π_n).

The model evolves by all agents making a contribution each period. At the end of each periods evidence is updated. By making contributions to a cluster evidence increases. The process is self-reinforcing. Contributions to a cluster increase the available evidence which in turn makes it more likely that agents will contribute to the same cluster next period.

In making his decision to which cluster to contribute a scientist looks at the available evidence. By evidence we understand the accumulated knowledge in the form of journals, textbooks and the like. Evidence is produced by scientists making contributions. We assume that all contributions produce evidence for just one single cluster. The quality of contributions is assumed to be homogeneous. Using these assumptions we avoid the task of having to judge the quality of contributions. We can measure evidence as the weighted sum of contributions, where evidence produced within a period equals the number of contributions within this particular period.[2] The evidence for cluster j at time t is given by $E_j(t)$ where $K_j(t)$ is the number of contributions to cluster j in period t and $d \in [0,1)$.

$$E_j(t) = K_j(t) + dK_j(t-1) + d^2 K_j(t-2) + \ldots = \sum_{t=0}^{T} d^t K_{T-t} \tag{2}$$

The number of contributions to cluster j in period t are given by $K_j(t) = \sum_{n=1}^{N} a_{nj}$ where a_{nj}=1 if $j \in \underset{k \in \{1,2,\ldots,J\}}{\arg\max}\, \pi_{nk}$ and 0 else.

The basic model is a nonlinear Polya process with the probability that a new contribution is made to a specific cluster being a function of the contributions already made to that cluster. As previous choices matter and increase the probability that a contribution will be made to a cluster, this process is path-dependent and exhibits positive feedback. We are interested in the structure that emerges during this process, where by structure we understand the proportion of agents working within each cluster. As has been shown by Arthur, Ermoliev, and Kaniovski (1983, 1984, 1987), the structure, which in our model is a vector of proportions, tends to a limit random vector. Our model reaches a stable pattern if $E_{t+1}=E_t$. Since our agents face the same evidence each period, they make the same choice each period and the distribution of agents across clusters stays constant.

Knowing that a stable pattern emerges our next question is concerned with the size of the clusters. Do all clusters have rougly the same size or does one cluster become dominant? Assuming that preferences are drawn from a [0,1] uniform distribution, it is clear that for $c=0$ all clusters are roughly equal in size. In this situation agents only care about their intrinsic preferences, there is no premium on compatibility (no trust) and hence there is no positive feedback. Since preferences are uniformly distributed all clusters are of roughly equal size. However, as soon as $c>0$ increasing returns kick in. Clusters with high evidence attract more contributions and grow up to a certain point. This is visualized in figure 1, showing three runs of a simulation with $N=100$ agents and $J=5$ clusters. The figure depicts the size of each cluster, measured as the share of agents contributing to the cluster, on the vertical axes (formally: $\frac{K_j(t)}{N}$). The horizontal axes measures time.

---

[2] By using the sum of contributions as a proxy for evidence it would also be possible to relate the model to scientometric data.

c=0	c=0.75	c=1.0
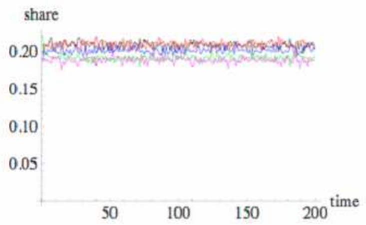	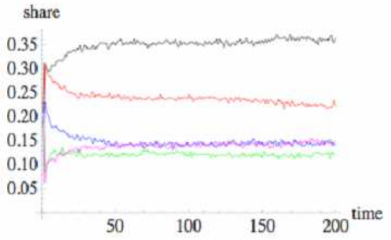	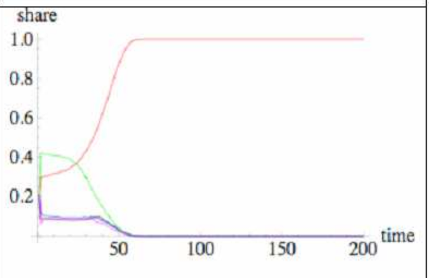

We see that the second and third simulations show sensitivity to initial fluctuations. Due to the initial fluctuations a cluster gets selected and starts growing faster than the other clusters. After about 100 time periods a stable structure emerged at which the size of the dominant cluster (and all other clusters) stays constant. A large cluster can be understood as the existence of high consensus and low disagreement, or much specialisation and low diversity. The maximum size of the dominant cluster increases with c, where if c increases above some threshold the dominant cluster gets 100%, i.e. there is absolutely no disagreement (no diversity) within the particular school of thought.

The resulting process exhibits several features of Arthur's increasing returns model (Arthur 1989). We cannot predict in advance which cluster will get dominant, but we can predict that one single cluster will get dominant; in Arthur's terms the process is *predictable*. The process is *nonergodic*, since small differences at the beginning (the distribution of preferences and initial evidence) are not averaged out over time, instead, they are responsible for selecting the dominant cluster. Having reached a stable state the size of all clusters stays constant, and the dominant clusters stays dominant. Consequently, the process is *inflexible*, i.e. there is no change from within the system.

In order to make the model more interesting and realistic we add some refinements. First, we introduce heterogeneity of agents. Agents are heterogeneous in how they react to available evidence, that is they differ in their c-parameter. Some agents have a high value of c_n indicating that they are followers rather than mavericks, who have a low value of c_n. Some agents are mavericks, and some are quite conservative, but most agents are somewhere inbetween. We reflect this fact by drawing c_n from a normal distribution with mean μ and standard deviation σ.
As a second modification we introduce sunk costs. Agents have invested a lot of time and money getting to know the methods used within the cluster. For them, it is not possible to appropriate the benefits from that investment if they switch to another cluster. The longer an agent has worked within a cluster, the less likely she is to change since her standing, reputation and accomplishments all depend on the correctness of the cluster. Let t_{jn} denote the number of contributions agent n made to cluster j and let $p_{jn} = \bar{p}_{jn}(1+t_{jn})^{1/\eta}$ where $\eta > 1$ is a parameter that reflects the responsiveness to sunk costs. Let p_n be the column vector with elements $(p_{1n}, p_{2n,\ldots,} p_{Jn})$. Now the elements of the matrix $P=(p_1, p_{2,\ldots,} p_N)$ are not fixed but change as agents make their contributions. The likelyhoods of pursuit are given by $\pi_n = p_n + c_n \hat{E}$. In the literature this is known as „hardening of positions“ where as time passes agents put more weight on their own opinion and less weight on the opinion of others (e.g. p. 4, Hegselmann and Krause 2002). It can also be understood as a process of dissonance reduction where agents adjust their preferences in order to reduce the discrepancy between their preferences and choices.

Third, our agents do not live forever. Each period their age increases by one unit. Once they have reached a certain age, randomly drawn from a uniform [50,100] distribution they die and get replaced by a new agent with age drawn from a uniform [20,50] distribution. The new agent faces the novice/expert problem, on which more will be said below.

Now let us clarify our computational model. The underlying model ontology is sparse, the model consists of N agents (indexed by subscript 1 to N) and J clusters (indexed by subscript 1 to J). The state of agent n is fully specified by the tuple (a_n,t_n,c_n), where a_n denotes age, $t_n=(t_{1n}, t_{2n,\dots,} t_{Jn})$, and c_n is the agent's c-factor, as explained above. The static state of our community of epistemic agents is fully specified by the tuple S and the vector E.

$$S=\{(a_n,t_n,c_n):1\leq n\leq N\}$$

$$E'=(E_1, E_2,\dots, E_J)$$

The evolution of the state of our scientific community, $S_t \rightarrow S_{t+1}$, is determined by simple rules which are repeatedly applied each period.

Step 1: If an agent reaches his maximum age he dies and gets replaced by a young agent. The young agent is assumed to be at the start of his career without having detailed knowledge about the available evidence. He faces the novice/expert problem which is modeled as a standard Polya process. The probability that the young agent makes his first contribution to cluster j is given by $\hat{E}_j$.

Step 2a: With a small probability ε the agent makes a contribution to a randomly chosen (but existing) cluster. Call ε the probability of mutation. Higher ε means that agents are more likley to try something new.

Step 2b: With probability (1-ε) the agent makes a contribution to the cluster in which she has the highest likelyhood of pursuit.

Step 3: Evidence is updated according to equation (2).

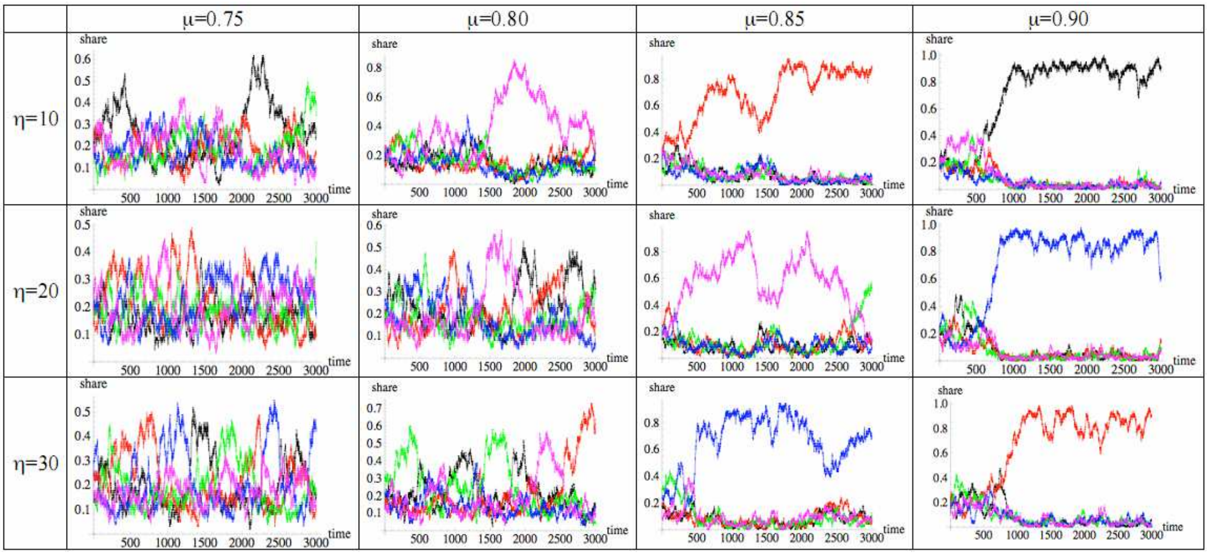

We did some comparative dynamics with the simulation and figure 2 shows simulation runs for varying η (from top to bottom) and varying μ (from left to right).[3] With increasing η the simulation shows more fluctuations and clusters stay dominant for a shorter period of time. If we increase η, p_n increases slower as agents make their contributions. This is a consequence

---

[3] If not indicated otherwise, parameters are set as follows: N=100, J=5, d=0.8, ε=0.025, initial preferences are drawn from a [0,1] uniform distribution, values for c are drawn from a Gaussian distribution with mean μ and standard deviation σ=0.1. The simulation code is avaible from the authors upon request.

of sunk costs increasing slower. Differences in preferences are smaller, preferences are smaller relative to trust times evidence ($c_n\hat{E}$), and agents are more likely to change clusters.

Figure 2 (from left to right) shows what happens if we change μ. Remember that μ determines how strong agents react to available evidence. From the figure we see that if we increase μ the dominant cluster gets larger, indicating a positive relationship between μ and the size of the dominant cluster. For large values of μ a single cluster gets a dominant position at 80-90%. „Mutating“ agents prevent the cluster from getting 100%. Since it is very likely that the cluster stays dominant we call this a *quasi lock-in*.[4] Our findings can be summarized as follows:

(1) If there is high trust (large μ) more weight is put on evidence. As agents put more weight on evidence larger clusters emerge, i.e. there is consensus among a larger group of agents. Scientists are more likely to specialise.

(2) If agents are more hard-wired (lower η) preferences p_n increase faster and agents are more conservative. This implies that older agents put a relatively low weight on evidence and are less likely to change clusters. Periods of consensus among large groups of agents are shorter, i.e. large clusters break down faster. Scientists are more likely to diversify.

(3) For large values of μ we get a quasi lock-in. This means that one cluster gets selected and grows up to about 80-90%, i.e. there is consensus among the vast majority of agents. More general, the size of the largest cluster increases with μ. Due to positive feedback effects large clusters are likely to stay large, and the probability of the largest cluster breaking down being inversely related to its size.

4. Conclusion

The simultaneous existence of multiple empirically adequate and coherent clusters reflects the *underdetermination* of theory by data. This evidence for each cluster will be *empirically adequate* because contributions to a cluster provide evidence for that cluster and the core of the research programme doesn’t come into question (cf. Lakatos). The basic assumptions of the contributions to a cluster will be *coherent* (because belonging to a cluster requires compatibility). This gives all clusters the potential to grow into a dominant research programme that will be coherent and empirically adequate and thus provide some kind of useful knowledge. However, what our model entails is that this success does not guarantee that the dominant cluster is also the *best* one. In other words, success does not entail truth.[5] Furthermore, the path-dependence that makes the model sticky such that conflicting new evidence is not immediately reflected in abandonment of a cluster adds a *historical* dimension to the analysis of science. The importance attributed in this model to network externalities reflects an incorporation of the *social* dimension of knowledge. The clusters described evolve without a definite end-point in sight and with the possibility that dominant clusters might one day be replaced (cf. pessimistic meta-induction). As such, the pattern resulting from our model can account for the emergence of consensus and the success of science, while also explaining the diversity of views and the possibility of pessimistic meta-induction.

---

[4] If we would set the probability of mutation to zero (ε=0) one cluster would lock-in at 100% (for large values of μ). With ε>0 it is still possible but very unlikely that the dominant cluster breaks down.

[5] *“The usual policy of letting the superior technology reveal itself in the outcome that dominates is appropriate in the constant and diminishing-returns cases. But in the increasing returns case laissez-faire gives no guarantee that the ‘superior’ technology (in the long-run sense) will be the one that survives.”* (Arthur 1989, 127)

Our model does justice to the idea that scientists are *rational*. It is, however, rational to be social because no scientist, however genial or productive, can collect all evidence on his own. This brings them to be nodes in a larger network with rules of its own. Rationality, provided c is large, leads to the emergence of a broad consensus. However, an important implication of this model is that the consensus does not necessarily entail that the cluster on which a consensus has developed is also the *right* one, nor does it imply that this consensus will last. The model has room for rationality, consensus and success. But individual rationality doesn't necessarily carry over to the whole, consensus is never absolute and success is no guarantee for truth. However, scientific activity *needs* to be mediated through the social in order to achieve a spectacular gain through specialisation and distribution of labour. But the dynamics of networks and the increasing returns by which these scientific networks are characterised are not truth-tropic. As such, dissensus is an essential part of any scientific community and science policy should consider taking on the role of antitrust agencies. Our model not only supports these points in principle, but should, through field-specific estimation of the relevant parameters and calibrated against historical episodes, enable policymakers to assess the results of policy decisions. It is hoped that the mechanism described can lead to better institutional design.

In subsequent work we will further refine the basic model. So far, we have assumed that the c-factors are all drawn from the same distribution. In a more interesting variant of the model there would be a specific c-distribution for each cluster. Mean μ and standard deviation σ of c will be a specific feature of the cluster, with high μ indicating a high potential for growth.
We can explain how consensus emerges. The breakdown of the dominant cluster is triggered by mutation. That is, by allowing for mutation we are able to explain the dissolution of consensus. (A high probability of mutation could by explained by a large number of anomalies in the dominant cluster.) Another way to get some more interesting dynamics is the idea of *unification of clusters*. By this we mean that two clusters merge together into one new cluster. In the history of science this happened with Newton unifying Galileo and Kepler, or the neoclassical-Keynesian synthesis. Also, we plan to introduce the possibility that clusters split up and new clusters get born.

Science is a distributed effort. This means that scientists, for a large part of their research, are building on knowledge they have not acquired first-hand but through education, reading and reports from colleagues. Division of labour enables scientists to be extremely more productive. They can specialise in a certain subfield, focus on specific problems and leave experiments unperformed. While in classic epistemology the focus seems to lie more with the knowledge a scientist acquires first-hand, *social* epistemology has the intention of dealing with these broader topics of trust and testimony between scientists. While classic epistemology focuses on the evidence for the truth of p, this paper draws attention to the fact that at the moment when epistemology goes social, *network externalities* come to have a large influence on the acceptance of p. The distributed character of knowledge production means that network externalities have an undeniable role to play in the assessment of knowledge claims. This paper has constructed a model which illustrates the dramatic effects of network externalities on a community of rational scientists working in a distributed fashion.

REFERENCES

Arthur, W. Brian (1989), Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, Vol. 99, pp. 116-131.

Arthur, W. Brian, Yu. M. Ermoliev and Yu. M. Kaniovski (1983), A generalized urn problem and ist applications, *Cybernetics*. Vol. 16, pp. 61-71.

Arthur, W. Brian, Yu. M. Ermoliev and Yu. M. Kaniovski (1984), Strong laws for a class of path-dependent stochastic processes, with applications. In: Arkin, Shiryayev and Wets (eds), *Proceedings of the Conference on Stochastic Optimization, Kiev 1984,* Lecture Notes in Control and Information Sciences, Springer.

Arthur, W. Brian, Yu. M. Ermoliev and Yu. M. Kaniovski (1987), Path-dependent processes and the emergence of macro-structure, *European Journal of Operational Research*. Vol. 30, pp. 294-303.

Goldman, Alvin (2001), Experts: which ones should you trust?, *Philosophy and Phenomenological Research*, Vol. 63, No. 1, pp. 85-110.

Hardwig, John (1985), Epistemic dependence, *Journal of Philosophy,* 82: 335-349.

Hardwig, John (1991), The role of trust in knowledge, *Journal of Philosophy*, 88: 693-708.

Hegselmann, Rainer and Ulrich Krause (2002), Opinion dynamics and bounded confidence: models, analysis and Simulation, *Journal of Artificial Societies and Social Simulation*. Vol. 5, <http://jasss.soc.surrey.ac.uk/5/3/2.html>.

Kitcher, Philip (1990), The division of cognitive labor, *Journal of Philosophy*, Vol. 87(1), pp.5-22.

Kuhn, Thomas (1962). The structure of scientific revolutions. Chicago, Chicago University Press.

Laudan,Larry (1981), A confutation of convergent realism, *Philosophy of Science,* Vol. 48(1), pp. 19-49.

Laudan, Larry (1984), *Science and values,* London:University of California Press.

Mirowski, Philip (2004), The scientific dimensions of social knowledge and their distant echoes in 20th-century American philosophy of science, *Studies in History and Philosophy of Science*, Vol. 35(2), 283-326.

Putnam, Hilary (1975), *Mathematics, matter and method.* Cambridge: Cambridge University Press.

Zamorra-Bonilla, Jesus (1999), The elementary economics of scientific consensus, Theoria, 36, 461-488.