Routledge
Taylor & Francis Group

The division of labour in science: the tradeoff between specialisation and diversity

Rogier De Langhe*

Centre for Logic and Philosophy of Science, Ghent University, Blandijnberg 2, Room 2.26, 9000 Ghent, Belgium

Economics is a typical resource for social epistemology and the division of labour is a common theme for economics. As such it should come as no surprise that the present paper turns to economics to formulate a view on the dynamics of scientific communities, with precursors such as Kitcher (1990), Goldman and Shaked (1991) and Hull (1988). But although the approach is similar to theirs, the view defended is different. Mäki (2005) points out that the lessons philosophers draw from economics can go either way depending on the model chosen. Thus, the aims of this paper are (1) to illustrate this flexibility by proposing an alternative model which assumes increasing returns to adoption in science rather than the decreasing returns present in the aforementioned contributions; and (2) to outline the implications of this view for scientific pluralism and institutional design.

Keywords: economic epistemology; division of labour; increasing returns; network industries; scientific pluralism; institutional design

1 Introduction

The subject of this paper is the division of labour in science. Should scientists diversify or specialise? The larger a consensus within a group of scholars, the more scholars can reap the benefits from dividing labour and the specialisation it allows; the more dissensus, the richer the diversity of views. Both specialisation and diversity have been attributed epistemically beneficial features. However, they cannot simply be maximised simultaneously because they are inversely related; more diversity means fewer possibilities for specialisation and vice versa. As such, a *tradeoff* obtains between the two. This gives rise to a number of related questions. Is there an optimal distribution we should aim for in science? Is it specific for a discipline, domain, group, question, ... or science in general? Does it change throughout history and if so, what are its historical and social determinants? *Can* we change an existing balance, *should* we do it and *how* can it be done? And who are *we* anyway?

Any group, academic or other, occupies at any time a certain point in between these two extremes of specialisation and diversity. To determine whether that point is also the *right* one, an answer is needed to the question of how to distribute labour. This means that the relevance of this question is both theoretical and practical. The broader implications lie, theoretically, with scientific pluralism (i.e. the normative endorsement of a plurality of views in science). An answer to the question how to distribute labour in science can be seen as a measure for the desirability of pluralism. Also, analysing pluralism from the

*Email: Rogier.DeLanghe@Ugent.be

perspective of the tradeoff between specialisation and diversity provides more clarity on the *kind* of pluralism that is preferable. The practical significance is that it can yield guidelines for institutional design and arguments to be used by (science) policymakers.

The paper is organised as follows. Section 2 explains *what* this tradeoff is and how it relates to science. Section 3 introduces increasing returns and relates them to scientific knowledge. Section 4 explores the analogy between science and network industries; section 5 then offers an application of this analogy to the economics discipline. The view of science as a network industry is used in section 6 to formulate an answer to the question which position on the tradeoff *should* be preferred. Finally, section 7 examines how the desiderata from the previous section *can* be achieved and section 8 concludes.

2   Specialisation and diversity

The question of the division of labour in science has been characterised by Helen Longino 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' (Longino 2006). In other words, the question whether science should *specialise* or *diversify*. The benefits of diversity have been widely recognised, for example by Philip Kitcher, who states that 'we sometimes want to maintain cognitive diversity even in instances where it would be reasonable for all to agree that one of two theories was inferior to its rival, and we may be grateful to the stubborn minority who continue to advocate problematic ideas' (Kitcher 1990, p. 16). Similarly, Paul Feyerabend wrote 'that progress can be brought about only by the active interaction of different "theories"' (Feyerabend 1970, p. 203). More generally, the merits of diversity spring from the fact that alternatives provide each other with valuable criticism (Mill, Popper, Feyerabend, Longino) and enable a division of labour (Hull 1988; Kitcher 1990; Goldman 1991). In the economy, the presence of a diversity of firms, products and preferences is said to bring about competition, higher productivity and a greater collective benefit.

On the other hand, scholars like Thomas Kuhn hold that the main emphasis should lie with specialisation: 'At least for the scientific community as a whole, work within a well-defined and deeply ingrained tradition seems more productive of tradition-shattering novelties than work in which no similarly convergent standards are involved.' (Kuhn 1977, p. 234). And of course the economic benefits associated with specialisation have become received wisdom ever since Adam Smith's tale of the pin factory.[1] According to Smith, specialisation is beneficial because it prompts a strict definition of tasks, leads to the development of more specific tools and allows workers to achieve a higher level of personal skill.

In sum, both specialisation and diversity have been attributed epistemically benign features. But, perhaps disconcertingly, they are inversely related to each other: more diversity means fewer possibilities for specialisation and vice versa. However, there is an obvious sense in which specialisation and diversity are not at odds at all. For example at an assembly line, a paradigm example of specialisation, there is a large amount of diversity because the workers perform different, complementary tasks. Specialisation here boosts diversity rather than impedes it. The reason why the tradeoff fails to obtain here is because two levels are being confused: the level of the standard which enables the division of labour (the assembly line) and the higher level which sees *across* such standards (cf. infra for more on standards). Within the standard, at the level of the assembly line itself, there is indeed diversity; but seen from the higher level, there is only one standard at work. It is exactly because everyone shares the same standard that the division of labour at the

assembly line is so efficient. If different standards were to operate simultaneously, there would be more diversity but productivity would suffer. So both specialisation and diversity must explicitly be situated at the inter-standard level in order to make sense of the tradeoff as it was introduced. A discipline is then specialised when the community's research efforts are concentrated on a small part of the available standards; it is diverse when research effort is equally distributed among standards.

Given that both are said to be epistemically beneficial and cannot be maximised simultaneously, a balance has to be struck. Somehow specialisation and diversity need to be traded off against each other. This problematic relationship has not remained unnoticed. Thomas Kuhn calls it 'the essential tension'[2] between traditionalism and iconoclasm; Paul Feyerabend states that the interplay between what he calls 'the principle of tenacity' and 'the principle of proliferation' (roughly equivalent to what I call specialisation and diversity) is an essential feature of the actual development of science.[3] Larry Laudan (1984) pointed out that no satisfactory account has been developed which simultaneously explains consensus and divergence in science.[4] But even going back to Adam Smith leads us to a corollary of this tension,[5] namely the tension between the already mentioned pin factory and that other famous concept of his: the invisible hand. Smith uses the example of the pin factory to argue that the amount of specialisation is limited by the extent of the market; the bigger the market, the more specialisation and the more economies of scale. The ideal factory size is thus one that dominates the entire market, i.e. a monopoly. On the other hand, the invisible hand which is necessary for optimal allocation can only function properly in a world of perfect competition. As such, the pin factory and the invisible hand are, strictly speaking, inconsistent.[6] Actually they might be compatible under special conditions, but the increasing and decreasing returns, for which they are a metaphor, are not. Increasing returns set in motion a virtuous circle that make it tend toward monopoly. On the other hand, in a context of decreasing returns, optimal size is usually smaller than the entire market, leaving room for alternatives. Because of the relation between the assumption of increasing and decreasing returns on the one hand and the tension between specialisation and diversity on the other hand, the assumption about returns to adoption will play a key role in this paper.

3   Knowledge and increasing returns

Increasing returns are best known in the form of increasing returns to scale, a part of standard economic theory. As a company gets bigger it can realise a higher level of efficiency. These gains are linear and affect only the company involved. Returns to scale are a common occurrence in the literature on scientific development. For example the theory of the cycle of credibility, in Latour and Woolgar (1986), which demonstrates the mechanism of a virtuous cycle of obtaining credit, transforming it to more resources for new research and hence obtaining even more credit. These benefits are linear and restricted to the individual. By contrast, networked increasing returns are exponential and the gains are spread throughout the entire network. A telephone network is a good example. A telephone is worthless without other people to call; its value increases as more people get one. As these networks grow, their value grows exponentially and for all adopters; one extra adopter enables two calls, two extra gives four additional connections, etc. Crucially, the value of the network depends on the number of agents adopting a certain standard.

Note that the level of analysis is here the activity *between* different communities adopting different standards. Within a certain community there might well be free competition among scientists, but across communities this competition is thwarted by transaction costs involved in switching standards. The point is that if one takes these

R. De Langhe

standards seriously, the problem of the division of labour changes from the dynamics within a consensual community (as in Kitcher) to the dynamics between standards competing for adoption. From an institutional design perspective, this is the relevant level because it's here that science policy goals can be pursued.

In economics, the main difference between the pin factory and the invisible hand is that the former is driven by increasing returns and the latter by decreasing returns. The bigger the pin factory, the more specialisation is possible and the higher the profit. Hence there is no optimal company size; the best size is determined by the extent of the market. On the other hand, the invisible hand relies on competitive equilibrium, with an optimal company size due to decreasing returns. Although there are plenty of instances in the economy characterised by increasing returns, such as railways, telecoms, energy,[7] a lot of contributions in economic theory have traditionally presupposed decreasing returns. This is perhaps because decreasing returns make it possible to find a single, optimal equilibrium and enable the use of clean and powerful mathematical tools. Increasing returns, on the other hand, entail multiple and possibly suboptimal equilibria, path-dependence and sticky prices, and cause economies to lock into inefficient technologies. In other words, assuming decreasing returns is 'the path of least mathematical resistance' (Krugman 1994, p. 4).

Although economics is currently undergoing an 'increasing returns revolution', philosophers of science drawing on economics have tended to follow the most accepted strands of traditional theory and their assumptions, including the decreasing returns one: models of scientific activity such as Kitcher (1990), Goldman and Shaked (1991) and Hull (1988) presuppose that scientific activity is characterised by decreasing returns to adoption. Is this based on an explicit argument for the prevalence of decreasing returns in science or have they just turned to economics in search of tools and picked those tools which were best elaborated? Whether increasing or decreasing returns are prevalent in science is ultimately an empirical matter, but in the meantime it makes sense to develop an alternative view that questions the implicit assumptions of those previous models and explores their consequences. As Mäki (2005) has pointed out, the models in economics from which social epistemology can draw can be used to make different and incompatible points about science. This means Kitcher's lesson can be turned upside down. To this effect, I will point out how social aspects of science can cause specialisation rather than diversification depending on a different assumption about returns to adoption. In particular, in Kitcher (1990), the number of adopters negatively affects the number of future adopters because it limits agents' chances to be the first to come up with the solution. But if science is seen as a network activity (distributed), the stronger – i.e. the larger – the network, the more agents will want to be compatible with it. Hence the number of adopters *positively* affects the number of future adopters.

So what would such an alternative model look like? The value of the (product of the) network does not primarily depend on scarcity, what it is made of or who made it, but on how many people use it. Consequently, seeing scientific communities as networks means they increase exponentially in strength every time a node is added to the network. As such, the addition of one scientist to a scientific community wouldn't decrease the other scientists' prospects of success (as in Kitcher) but would increase the prospect of success for the community exponentially. It would be a model of scientific activity under increasing returns to adoption.

In a series of papers, Brian Arthur has tried to find a formal way of capturing the dynamics of systems exhibiting such increasing returns to adoption. Arthur (1989) points out that many increasing-return problems fit a general nonlinear probability schema. In Arthur (1994), the general idea is laid out as follows:

Figure 1.   Increasing returns adoption: a random walk with absorbing barriers.
Source: Adopted from Arthur (1989), p. 120.

> It can be pictured by imagining a table to which balls are added one at a time; they can be of several possible colors – white, red, green or blue. The colour of the ball to be added next is unknown, but the probability of a given color depends on the current proportions of colors on the table. If an increasing proportion of balls of a given color increases the probability of adding another ball of the same color, the system can demonstrate positive feedback. (Arthur 1994, p. 6)

The result of such a system is a random walk with absorbing barriers: the ratio of colours varies randomly but if the ratio of any colour reaches a certain level (the barrier) it will never get under that level afterwards. In other words, the outcome is locked in.

A paradigm for such models has been set in Arthur (1989), in which a simple example is provided in which two technologies, A and B, compete for adoption. There are also two types of agents: R-types choose A and S-types choose B. Agents are thrown into the game with a 0.5 probability of being an R-type or an S-type. This probability changes once the amount of adopters of A or B reaches a certain threshold, after which all new agents adopt to the dominant technology irrespective of their being an R- or S-type. The resulting dynamic is one of a random walk with absorbing barriers. Interestingly, it is certain that agents will eventually lock into one of the two technologies, but which one is impossible to predict; it is different every time the game is played. More generally, Brian Arthur notes that the characteristics of the dynamics of this game will differ in crucial respects from the dynamics of games assuming decreasing returns: it will exhibit multiple equilibria, non-predictability, potential inefficiency, inflexibility and non-ergodicity.

4   Science as a network industry

The main focus of the present paper lies in the tradeoff between specialisation and diversity and its implications for pluralism. For a detailed formal model applying increasing returns to scientific activity, see De Langhe and Greiff (in press). In this paper the emphasis is on the development of the analogy between scientific activity and network industries.

Hans-Werner Gottinger gives the following definition of a network industry:

> Network industries can be defined as those where the firm or its product consists of many interconnected nodes, where a node is a unit of the firm or its product, and where the connections among the nodes define the character of commerce in the industry. (Gottinger 2003, p. 1)

In a similar sense, I take a scientific community to consist of a network of nodes. The nodes are individual contributions (a paper, a book, introduction of a new idea). A necessary condition for these nodes to form a network (a scientific community) is that there is a *standard* which makes the nodes compatible with each other. Each time a scientist makes a contribution, the contribution falls under a certain standard and the scientist can be said to 'adopt' to that standard at that point. Adopters to a certain standard form a community and different communities based on different standards then compete at the level of the discipline. It is important to note that the very possibility of a division of labour in science crucially depends on the presence of such a standard. Without a shared standard within which different actors can perform different tasks but still work on a common project it would be impossible for scientists to coordinate their research efforts. Standards are a condition of possibility for the aggregation of individual contributions and for accumulation of the results over time. As such the topic of division of labour cannot be addressed without at least implicitly presupposing the presence of a standard. This paper aims to make this presupposition explicit and explore its consequences by making the analogy between the dynamics of scientific activity within a discipline and the dynamics of standards competing for adoption.

There are a number of good prima facie reasons for exploring this analogy. First, as emphasised in recent contributions such as Giere and Moffatt (2003) and Hutchins (1995), scientific activity is a distributed effort, much like the individual nodes forming the network featured in the definition of a network industry. Secondly, scientists seem to cluster their ideas around certain basic concepts and key theoretical assumptions, for example those called 'paradigms' by Kuhn (1962); these act as *standards* with which scientists comply. Thirdly, information is characterised by increasing returns. This has important effects on companies that produce them, for those companies function differently from other, mainly industrial, companies. Fourthly, a bigger network also means more feedback, more conferences to go to, more journals to publish in, more jobs to apply for, etc. Networks do seem to matter in the academic world. A fifth factor is that the incentive structures at universities tend to reward research that has been successful in the past. Those rewards often come in the form of more research means, enabling the successful researchers to be even more successful in the future. This has even grown with the increasing importance of bibliometric methods, which has made joining a larger network even more compelling. A simple example can illustrate this powerful feature of networks. A discipline has two research communities (RC). RC1 has three members, RC2 has four. Now let's say all members publish a paper and cite all other members in that paper. In RC1 this will give each member six citations, in RC2 the extra member *doubles* the number of citations (12). This provides a strong incentive for a newcomer to join RC2, which would in effect make the discrepancy between RC1 and RC2 even stronger (20 vs 6). As such, the larger the network with which a scientist's contribution is compatible, the larger its impact.

Using Arthur's paradigm as a starting point, the following analogy between science and network industries is 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 to cluster x depends on the share of x in the discipline. This results in increasing returns to adoption because the more people join a cluster, the bigger the odds that a newly born scientist will join that cluster. An idea by Alvin Goldman helps to make this step intuitive. Upon entry, the newborn scientist finds himself in what Goldman called the novice/expert-problem.

> 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)

I maintain that the novice is incapable of making his choice on the basis of scientific content. By definition, the novice has not yet spent any energy in getting to know the field, so he cannot rely on his knowledge of the different clusters, he cannot assess the reliability of different experts and he has no perfect knowledge of the state of the discipline which would enable him to compare the available evidence for different clusters or to know which network has the most adopters. Instead, what I propose is simply that the probability that the novice contributes to network x is equal to the share of x in the discipline. This can be interpreted as the idea that the novice is likely to join the cluster 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 is that we require agents to be perfectly informed about the state of the discipline or have some prior knowledge about its contents (which would make novices not so novice after all).

Once a scholar is in the field, he will get to know the available evidence and arguments. This will reinforce the mechanism that larger clusters increase the odds of a scientist joining, because the more adopters a cluster has, the better the quality will be of the arguments and evidence produced in it. Moreover, as a scholar stays longer in the cluster, sunk costs increase, making change of network more costly. So in the model proposed, a scientist shifting between networks is possible but becomes increasingly unlikely as the scientist ages.

Then what does it mean to join a network? Each network consists of nodes (contributions) that are held together by a certain *standard*. The only contributions that can be added to the network are those compatible with that standard. As for scientific communities, a standard specifies which basic concepts to use, what theoretical assumptions to make, what the important problems are and what counts as a good solution. Note that the way a standard in science is characterised here is analogous to concepts such as 'paradigm', 'research programme', 'schools of thought', etc. The analogy between science and network industries then means that a scientist working within a certain paradigm can be approached in the same way as a computer user adapting to an operating system (e.g. Apple, Windows).

An account based on increasing returns means that science is driven by compatibility. Why are scientists more likely to adapt to bigger networks? Because there are *network externalities*. The strength of the network rises exponentially with the number of adopters. The larger the network a scholar contributes to, the more people he will reach, the more impact it will have and the less costly his effort will have been. The bigger the network, the more gains from network externalities and the more attractive it is to join it. This makes sense, at the aggregate level, the intrapersonal level and the personal level. At the aggregate level, scientists are involved in a never-ending quest to increase the size of the network they adapt to,[8] refining their own network to resist acquisitions by others and discussing with others in order to find common ground and spot potential synergies. From this point of view, scientific research can be seen as performing due diligence; unification and 'economics imperialism' can then respectively be seen as mergers and acquisitions. This drive for compatibility provides an explanation for 'unification' as an ideal in science. At the interpersonal level, the drive for compatibility explains why scientists try to persuade each other. In a world like Kitcher's, there is a premium on having as little as possible other scientists pursuing the same path as yourself, because every new scientist means a new competitor (i.e. decreasing returns to adoption). Scientists travelling the

world to share their ideas and making efforts to persuade one another is a truly puzzling occurrence from this point of view. It does not make sense in an environment with decreasing returns to adoption. If a scientist believes he is on the right track, the last thing he will want to do is to inform others, because this increases the other's chance of winning the prize and decreases his. The assumption of increasing returns to adoption, however, can account for this otherwise irrational behaviour. Thirdly, on a psychological level, the drive for compatibility seems to be hard-wired in the human brain since people tend to avoid cognitive dissonance.

What can be said about the epistemic quality of the result of this process? At the individual level, scientists try to reap the benefits of network externalities. At the aggregate level, each network produces its set of coherent, empirically adequate evidence; however, the networks are more or less incompatible (or, in more philosophical terminology, incommensurable) in the same sense as Microsoft and Macintosh. On the one hand, it is understandable why incommensurability elicits such horror (for it implies serious limits to compatibility). However, a network perspective also injects this prospect with some hope, because, as Gottinger notes about network industries: 'compatibility should be conceptualized as a continuüm' (Gottinger 2003, p. 5). In networks, 'converters' can be developed between networks, e.g. a program to run Microsoft software on a Macintosh computer. These converters (cf. bridge principles) can work both ways, as could be said to be the case when unification in science takes place, but it might as well only work one way (as was the case with running Microsoft software on a Macintosh computer). In science, the equivalent of the development of a one-way converter might be the discovery of theoretical tools to reduce one body of theory to another.

How should different networks then be evaluated? Network industries tend to develop a dominant standard (or, applied to science, a dominant paradigm). What is its status? Scientific realism would hold that a successful cluster of ideas, successful in the sense of being coherent, empirically adequate and popular among scientists, is most likely to be closest to the truth. However, a model exhibiting positive feedback mechanisms like the Arthur model implies that a specialisation bias is inherent to its dynamics. In a system driven by a quest for compatibility, the system as a whole is likely to bring up a dominant standard without the necessity to assume that this is the result of any truth-tracking ability possessed by the dominant standard. Arthur's paradigm model described in section 3 yields a dominant technology every time it is played. However, even with fixed initial conditions, the technology dominating varies every single time the game is played, indicating that the popularity of the technology does not depend on any inherent qualities of the technology itself but rather is brought about by network externalities. As such, within this perspective success is no proxy for truth.

5 Economics as a network industry

By way of illustration, let me apply this view to the discipline of economics. Although there is an obvious circularity in assessing the economics discipline through a perspective based on one of its own models (cf. Mäki 2005, p. 218), I do believe it might offer a fruitful alternative perspective to counteract the influence of models assuming an invisible hand (and often, as a consequence, decreasing returns) without even being explicit about it.

Science as a network industry means that there is a significant specialisation bias. Can any signs of this be detected in economics? The dominance of 'neoclassical economics' has been lamented by so-called 'heterodox economists' and pleas for more pluralism in economics are widespread.[9] However, Colander Holt, and Rosser (2004) warn against

exaggerating the homogeneity of dominant strands in economics and emphasise the changes the discipline has gone through since the 1960s, away from what they call 'the holy trinity – rationality, selfishness, and equilibrium –' (Colander et al. 2004, p. 485). Sent (2006) even goes one step further and argues that economics has failed to become a unified discipline, despite several serious attempts.

However, there seems to be a rather widespread discontent with how economics is faring as a discipline, giving rise to e.g. the Post-Autistic Economics movement and economists calling themselves 'heterodox'. One of its main proponents, Tony Lawson, acknowledges the argument by Colander et al., but adds:

> But it remains the case that these and all other widely sanctioned examples of ongoing change, diversity, novelty, complexity, evolution and multi-dimensionality, etc., are occurring within the framework of formalistic modelling. The insistence on *mathematical–deductive* modelling prevails in all cases; the essential feature of the recent and current mainstream remains intact. (Lawson 2006, p. 491; emphasis added)

On this, Colander et al. seem to agree with Lawson:

> Our view is that the current elite are relatively open minded when it comes to new ideas, but quite closed minded when it comes to alternative methodologies. If it isn't modeled, it isn't economics, no matter how insightful. (Colander et al. 2004, p. 493)

Numerous authors have characterised the above view as a caricature of the discipline, but also the less outspoken variants point to an interesting puzzle: how can the relatively high rate of homogeneity in economic theory be reconciled with the low rate of decisive empirical information, to wit information that would provide substantive reasons to prefer certain approaches over other less empirically successful ones? The view of scientific activity as a network industry proposed in this paper suggests that this situation can be interpreted as follows: (parts of) the economics discipline has (been) locked into a 'mathematical-deductive modeling' standard. This would explain the repeated pleas for pluralism and is in line with the existence of shared basic concepts and theoretical assumptions within large parts of economics. The presence of a broadly accepted standard allows for extensive division of labour and hence results in strong specialisation within the standard and large differences in appeal compared to rival approaches based on less popular standards. The broad acceptance of a certain standard in economics can explain why economics textbooks tend to be a-historic, giving economics students little or no knowledge of rival approaches, while textbooks in other social sciences typically do contain information on the development of the discipline and possible alternative standards. Because of its large rate of adoption, the standard itself can be taken for granted and needs no explicit legitimation anymore. Given that the power of a network crucially depends on its size (it rises quasi-exponentially with the number of adopters), the broad acceptance of a certain way of doing economics gives rise to strongly integrated networks in which non-conformists are sanctioned.

> to get an article published in most of today's top rank economic journals, you must provide a mathematical model, even if it adds nothing to your verbal analysis. I have been at seminars where the presenter was asked after a few minutes, 'Where is your model?'. When he answered 'I have not got one as I do not need one, or cannot yet develop one, to consider my problem' the response was to turn off and figuratively, if not literally, to walk out. (Lipsey 2001, p. 184)

Simultaneously, the dominant standard is characterised by an abundance – almost to the point of redundancy – of slightly different contributions.

If design and manufacture have fixed costs, then there are too few products designed for networks or standards with few users, and too many for those with many users. This reinforces the winner-take-all characteristic of these markets. (Heal 1998, p. 12)

6   Scientific pluralism

Following the illustration of economics as a network industry in the previous section, it is now time for more general conclusions. How do increasing returns affect the tradeoff between specialisation and diversity? And what are the implications for pluralism? Pluralism involves a normative endorsement of diversity of views and as such it formulates an answer to what position on the tradeoff *should* be pursued. A tentative answer to this question can now be formulated.

The impact of network externalities on network industries produces a specialisation bias. By this I mean that a group of scholars (*ceteris paribus*, for example institutions which constitute the group have to be strong enough to support large groups) will be driven by the same forces that gives for example software companies a natural tendency towards lock-in/monopoly/winner-take-all. It provides scientists with an incentive to follow those strands which are already well established and to keep away from marginal research programmes. In other words, from this point of view the extra-scientific factors which Kitcher (1990) refers to[10] to avoid the CO–IR discrepancy are not enough to counter the effects of network externalities on scientific communities.

One of the lessons of Kitcher (1990) is that the impact of social factors on the division of labour in the sciences needn't be disruptive for the achievement of our epistemic goals, but might indeed be essential for it. If scientists were all rational, they might all choose the same research path, consisting of the one which has, based on the available evidence, the highest chance of success. Thanks to lesser motives such as the thirst for fame and fortune, scientists divide their labour across different research paths because a path's lower chance of success is compensated by the smaller amount of competitors. As such, his idea of what *should* be done becomes clear in this passage:

The very factors that are frequently thought of as interfering with the rational pursuit of science – the thirst for fame and fortune, for example – might actually play a constructive role in our community epistemic projects, enabling us, as a group, to do far better than we would have done had we behaved like independent epistemically rational individuals. Or, to draw the moral a bit differently, social institutions within science might take advantage of our personal foibles to channel our efforts toward community goals rather than toward the epistemic ends that we might set for ourselves as individuals. (Kitcher 1990, p. 16)

In Kitcher's model extra-scientific factors cause diversity and this diversity is then efficiently allocated by an invisible hand. Pluralism seems to be a matter of laissez-faire, with institutions merely to make sure this invisible hand can operate. On the other hand, a model incorporating increasing returns will follow a different dynamics, leading to different implications for pluralism and opposing laissez-faire.

To compare Kitcher's pluralism to the one proposed in this paper, I will position them in the classification developed by Van Bouwel (in press) which distinguishes three different views on scientific pluralism. Each one of them offers a different answer to the question what a good balance between specialisation and diversity should be and how to achieve that balance.

Consensual pluralism relies on the invisible hand to sort out different views. As such, the tension between specialisation and diversity disappears. No dissent is needed. Its attitude toward the tradeoff is one of laissez-faire. Groups will self-organise toward their optimal balance. An example is Kitcher (1990), claiming that diversity is necessary to

491   avoid a discrepancy between individual rationality and the collective optimum. The
492   related answer to what should be done is then a laissez-faire approach. Philosophers of
493   science in *turn* have – implicitly or explicitly – made use of the idea of the invisible hand
494   to explain how diversity and dissent can still be consistent with rationality and objectivity
495   in science.[11] This leads to pluralism without dissent.
496   *Agonistic pluralism* stresses the importance of dissent instead of relying on an invisible
497   *hand*. It suggests rules for active engagement such as those provided by Helen Longino.
498   The focus is not on finding the right balance between specialisation and diversity, but on
499   creating an adequate environment for fair discussion. Which concessions can be made
500   with respect to the tradeoff will be negotiated between communities of inquirers.
501   *Antagonistic pluralism* prefers to shift the balance toward specialisation, at the cost of
502   diversity. For example Kuhn sees innovation and revolution not as products of diversity,
503   but as the result of sustained specialisation.[12] As such he seems to be in favour of shifting
504   the balance towards more specialisation:

505
506   history strongly suggests that, though one can practice science [. . .] without a firm consensus,
507   this more flexible practice will not produce the pattern of rapid consequential scientific
508   advance [in which] development occurs from one consensus to another, and alternate
       approaches are not ordinarily in competition. (Kuhn 1977, p. 232)

509   Jouni-Matti Kuukkanen (2007) acknowledges that, especially in his later work, Kuhn
510   viewed the development of science as a strong tendency toward internal coherence
511   resulting in a plurality as a sum of exclusive monisms; a fragmentation of science.[13] As
512   such, antagonistic pluralism is the only standpoint which has an actual preference for a
513   certain position on the tradeoff, whereas consensual pluralism opts for laissez-faire and
514   agonistic pluralism sees it as a matter of negotiation.
515   The three views can be schematised as in Table 1.
516   While Kitcher's pluralism falls in the first category, the pluralism defended in this
517   paper falls into the third. But unlike Kuhn, the reasons to support this view are not
518   historical but come as a result of the view of science as a network industry. On this view,
519   specialisation is not so much a desideratum but a consequence of the assumption of
520   increasing returns itself. As such, there is no other option than to accept that science is
521   driven by an urge to specialise and a quest for compatibility. This means that I see no other
522   option than to defend an antagonistic pluralism, in which different, more or less
523   incompatible clusters of basic concepts and key theoretical assumptions strive for
524   dominance over a discipline, although this dominance is not stable and no guarantee for
525   superiority (respectively accounting for the existence of scientific revolutions and the
526   pessimistic meta-induction). The reason for this acquiescence with the specialisation bias
527   of science is that it produces enormous network benefits which would be lost with more
528   diversity. However, full specialisation should be avoided, for the same reasons as
529   governments enact antitrust policy in economic sectors where increasing returns play a
530   large role (cf. Microsoft vs. EU).
531   This view responds to consensual pluralism by opposing a laissez-faire attitude.
532   Systems exhibiting increasing returns do not tend to find an optimal equilibrium; rather

535   **Q10** Table 1.

Consensual	invisible hand	laissez-faire	'epistemic liberalism'
Agonistic	discussion	engagement/dissent	'epistemic democracy'
Antagonistic	specialisation	dominance dissidence/indifference	'epistemic hegemony'

they have the tendency to lock into suboptimal states. This means that arguments for laissez-faire, letting the 'market for ideas' sort itself out, are highly problematic in the context of network industries.

> The long-standing public policy concerns over network industries are not accidental, because those industries often embody two major and widely recognized forms of potential market failure: significant economies of scale – with the potential for monopoly – and externalities. (White 1999, p. 1)

In response to agonistic pluralism, the present view acknowledges the importance of dissent. Dissent is important in identifying possibilities for merger and for finding powerful gateways. However, dissent and discussion are already hard-wired into the idea of scientific research as due diligence and need no further encouragement from an institutional design perspective. Rather than stimulating discussion between different networks, the present view acknowledges the (albeit not total) differences between networks, their potential for specialisation and their singularity.

In the end, this version of antagonistic pluralism comes surprisingly close to democracy because it resembles the same structure: a few large competitors each having their own monist worldview, but limited in their claims to hegemony by democratic institutional structures (without which a lock-in would probably emerge, such as dictatorship).

7   The scope for institutional design

The previous section was concerned with how specialisation and diversity *should* be traded off. But ought implies can. *How* can a group's balance between specialisation and diversity be changed? Who or what decides? How can a good position be reached, or does it come by itself? From the perspective of science as a network industry a number of policy recommendations and other possibilities of action can be formulated.

Perhaps disconcertingly, the prospects for science policy in a context of increasing returns must remain modest because of the unpredictability of the measures taken.

> To the extent that small events determining the overall path always remain beneath the resolution of the economists' lens, accurate forecasting of an economy's future may be theoretically, not just practically, impossible. (Arthur 1994, p. 12)

However, according to Arthur, possibilities remain for policymakers:

> Steering an economy with positive feedbacks into the best of its many possible equilibrium states requires good fortune and good timing – a feel for the moments when beneficial change from one pattern to another is most possible. Theory can help identify these states and times, and it can guide policymakers in applying the right amount of effort (not too little but not too much) to dislodge locked-in structures. (ibid., p. 12)

The general challenge in a context of increasing returns seems to be to reap the benefits of specialisation while not eliminating diversity altogether. In fact, because of the presence of network industries in the economy, governments have already developed extensive regulation to do just that: anti-trust policy. 'The essence of anti-trust policy is to find ways of keeping these economics of scale or of standardisation while bringing in competition' (Heal 1998, p. 4). The implementation of anti-trust law is seen as self-evident in most market economies, so why not in science if science is seen as a network industry? The main goal should then be to avoid groups from locking in or at least locking in to suboptimal standards. For science, this could be achieved by increasing the possibility of scientists shifting between networks. One way to do this is to promote the creation of (what in network theory is called) gateways: 'gateways [. . .] create the possibility to have

different types of (incompatible) networks communicate and work together' (Gottinger 2003, pp. 4–5). However, the creation of bridge principles, reduction and unification seem to fall beyond the scope of science policy and are more a matter of genuine scientific labour. Indeed, science as due diligence implies that scientists are already on the case. A more feasible measure to reduce the odds of a lock-in could be to increase scientists' knowledge of alternatives. Returning to the case of economics discussed above, it is striking that history of economics has all but disappeared from economic curricula worldwide. Not only might this have contributed to the apparent lock-in in which it appears to find itself today, but also it might suggest a remedy. Greater awareness of the history of economic thought could increase liquidity in the economics discipline. Perhaps surprisingly, teaching the history of economic thought might thus turn out to be the academic equivalent of free trade. A second way to achieve this is to try and assuage the specialisation bias. For economics, but also for science in general, this could be done for example by reducing the importance attributed to bibliometric methods. As illustrated with an example in section 4, bibliometric methods increase the specialisation bias. Academic success is increasingly being defined in terms of bibliometric criteria. There is an ever-increasing incentive to publish in highly cited journals, which puts an extra premium on compatibility.

8 Conclusion

Economics is a typical resource for social epistemology and the division of labour is a common theme for economics. As such it should come as no surprise that the present paper has turned to economics to formulate a view on the dynamics of scientific communities, with precursors such as Kitcher (1990), Goldman and Shaked (1991) and Hull (1988). But although the approach is similar to theirs, the view defended in this paper was different. The aforementioned contributions assume decreasing returns to adoption. Mäki (2005) points out that the lessons we can draw from economics can often go either way depending on the model one chooses. Thus, the aims of this paper have been (1) to illustrate this flexibility by putting forth a model assuming increasing returns to adoption in science rather than decreasing returns, viewing scientific activity as a network industry; and (2) to outline the implications of this view for scientific pluralism and institutional design.

Notes

1. 'Each person, therefore, making a tenth part of the forty-eight thousand pins, might be considered as making four thousand eight hundred pins in a day. But if they all had wrought separately and independently, and without any of them having been educated to this particular business, they certainly could not each of them have made twenty, perhaps not one pin in a day' (Smith 1977 [1776]).
2. 'That is why I speak of an "*essential* tension" implicit in scientific research. [. . .] Very often the successful scientist must simultaneously display the characteristics of the traditionalist and of the iconoclast' (Kuhn 1977, p. 227; emphasis added).
3. 'I shall call the advice to select from a number of theories the one that promises to lead to the most fruitful results, and to stick to this one theory even if the actual difficulties it encounters are considerable, the *principle of tenacity*. [. . .] if change of paradigms is our aim, [. . .] we must be prepared to accept a *principle of proliferation*. [. . .] the interplay between tenacity and

proliferation [. . .] is [. . .] an essential feature of the actual development of science' (Feyerabend 1970, pp. 203–209; emphasis added).

4. '[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, p. 2; emphasis added).

5. This tension is noted in for example Heal (1998) and discussed at length in Warsh (2006).

6. The main argument for this inconsistency has been provided by Coase (1937).

7. Blinder, Canetti, Lebow, and Rudd (1998) estimate that 89% of firms are subject to constant or falling marginal costs.

8. It is the network that determines what the relevant questions are and what count as good solutions. Within a network, scientists will most likely be competitors, e.g. to be the first to solve a problem and win a prize. But at the network level, they have a shared interest in the expansion of the network relative to other networks because that means more people find the problem important.

9. E.g. 'Plea for a Pluralistic and Rigorous Economics', an advertisement in *American Economic Review*, calling for 'a new spirit of pluralism in economics, involving critical conversation and tolerant communication among different approaches'.

10. 'The very factors that are frequently thought of as interfering with the rational pursuit of science – the thirst for fame and fortune, for example – might actually play a constructive role in our community epistemic projects, enabling us, as a group, to do far better than we would have done had we behaved like independent epistemically rational individuals' (Kitcher 1990, p. 16).

11. This point is made by Ylikoski (1995, p. 35).

12. 'At least for the scientific community as a whole, work within a well-defined and deeply ingrained tradition seems more productive of tradition-shattering novelties than work in which no similarly convergent standards are involved. How can this be so? I think it is because no other sort of work is nearly so well suited to isolate for continued and concentrated attention those loci of trouble or causes of crisis upon whose recognition the most fundamental advances in basic science depend' (Kuhn 1977, p. 234).

13. 'Especially in his later writings, Kuhn argued more or less directly that scientists have attempted to increase coherence in the history of science. However, Kuhn did not make a claim that the total coherence of science has increased in the course of history, and moreover, it is hard to find any historical evidence for it. Kuhn held that scientists' attempts to improve the coherence of their theories paradoxically tend to decrease the total coherence of science, because that activity leads to the fragmentation of science. This kind of speciation or specialisation of scientific fields is how Kuhn came to understand scientific revolutions in the latter part of his career' (Kuukkanen 2007, p. 556).

References

Arthur, B. (1989), 'Competing Technologies, Increasing Returns, and Lock-in by Historical Events', *The Economic Journal*, 99, 116–131.

——— (1994), *Increasing Returns and Path Dependence in the Economy*, Michigan: University of Michigan Press.

Blinder, A., Canetti, E., Lebow, D., and Rudd, J.B. (1998), *Asking About Prices: A New Approach to Understanding Price Stickiness*, New York: Russell Sage Foundation.

Coase, R. (1937), 'The Nature of the Firm', *Economica*, 4(16), 386–405.

Colander, D., Holt, R., and Rosser, J. (2004), 'The Changing Face of Mainstream Economics', *Review of Political Economy*, 16(4), 485–499.

De Langhe, R., and Greiff, M. (in press), 'Standards and the Distribution of Cognitive Labour', *Logic Journal of the IGPL*.

Feyerabend, P. (1970), 'Consolations for the Specialist', in *Criticism and the Growth of Knowledge*, eds. I. Lakatos and M. Musgrave, Cambridge: Cambridge University Press.

Giere, R., and Moffatt, B. (2003), 'Distributed Cognition: Where the Cognitive and the Social Merge', *Social Studies of Science*, 33(2), 1–10.

Goldman, A. (2001), 'Experts: Which Ones Should You Trust?', *Philosophy and Phenomenological Research*, 63(1), 85–110.

Goldman, A., and Shaked, M. (1991), 'An Economic Model of Scientific Activity and Truth Acquisition', *Philosophical Studies*, 63, 31–55.

Gottinger, H. (2003), *Economies of Network Industries*, London and New York: Routledge.

Q8 Hardwig, J. (1985), 'Epistemic Dependence', *Journal of Philosophy*, 82(7), 335–349.

Q8 ——— (1991), 'The Role of Trust in Knowledge', *Journal of Philosophy*, 88(12), 693–708.

Heal, G. (1998), 'The Economics of Increasing Returns', PaineWebber Working Paper Series, No. 97-20. *http://ssrn.com/abstract=88528*.

Hull, D. (1988), *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*, Chicago: University of Chicago Press.

Hutchins, E. (1995), *Cognition in the Wild*, Cambridge, MA: MIT Press.

Kitcher, P. (1990), 'The Division of Cognitive Labor', *Journal of Philosophy*, 87, 5–22.

Krugman, P. (1994), *Rethinking International Trade*, Cambridge: MIT Press.

Kuhn, T. (1962), *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press.

——— (1977), *The Essential Tension*, Chicago and London: University of Chicago Press.

Kuukkanen, J. (2007), 'Kuhn, the Correspondence Theory of Truth and Coherentist Epistemology', *Studies in History and Philosophy of Science*, 38, 555–566.

Latour, B., and Woolgar, S. (1986), *Laboratory Life: The Construction of Scientific Facts*, Princeton: Princeton University Press.

Laudan, L. (1984), *Science and Values*, Berkeley: University of California Press.

Lawson, T. (2006), 'The Nature of Heterodox Economics', *Cambridge Journal of Economics*, 30, 483–505.

Lipsey, R. (2001), 'Success and Failures in the Transformation of Economics', *Journal of Economic Methodology*, 8(2), 169–202.

Longino, H. (2006), 'The Social Dimensions of Scientific Knowledge', *Stanford Encyclopedia of Philosophy*, *http://plato.stanford.edu/entries/scientific-knowledge-social/*.

Mäki, U. (2005), 'Economic Epistemology: Hopes and Horrors', *Episteme: A Journal of Social Epistemology*, 1(3), 211–222.

Sent, E. (2006), 'Pluralisms in Economics', in *Scientific Pluralism*, eds. S. Kellert, H. Longino and K. Waters, Minneapolis: University of Minnesota Press, pp. 80–101.

Smith, A. (1977 [1776]), *The Wealth of Nations*, Chicago: University of Chicago Press.

Van Bouwel, J. (in press), 'The Problem With(out) Consensus: The Scientific Consensus, Deliberation and Agonistic Pluralism', in *The Social Sciences and Democracy*, ed. J. Van Bouwel, London: Palgrave Macmillan.

Warsh, D. (2006), *Knowledge and the Wealth of Nations*, New York: W.W. Norton.

White, L. (1999), *U.S. Public Policy Toward Network Industries*, AEI-Brookings Joint Center for Q9 Regulatory Studies Books and Monographs.

Ylikoski, P. (1995), 'The Invisible Hand in Science', *Science Studies*, 8(2), 32–43.