

The division of cognitive labor: two missing dimensions of the debate

Baptiste Bedessem

Laboratoire PPL, Université Grenoble Alpes

Forthcoming in European Journal for Philosophy of Science (final draft)

Abstract

The question of the division of cognitive labor (DCL) has given rise to various models characterizing the way scientists should distribute their efforts. These models often consider the scientific community as a self-governed sphere constituted by rational agents making choices on the basis of fixed rules. Such models have recently been criticized for not taking into account the real mechanisms of science funding. Hence, the question of the utility of the DCL models in guiding science policy remains an open one. In this paper, we show that two unconsidered dimensions would have to be taken into account. First, DCL studies miss the existence of distinct levels of epistemic objectives organizing the research process. Indeed, the scientific field is structured as a system of hierarchical, interconnected practices which are defined both by their inherent purposes and by various superposed external functions. Second, I criticize the absence of ontological considerations, since the epistemological significance of pluralism is highly dependent on the nature of the object under study. Because of these missing dimensions, current DCL models might have a limited usefulness to identify good practices of research governance.

1-Introduction

How should we allocate resources to scientific research? This question, although old¹, is of growing interest among scientists and policy makers. A current set of concerns revolves around the mechanisms of project selection by funding agencies, the role that peer-reviewing should play in grant allocation, and the institutional conditions promoting the development of

¹See, for instance, Wilholt and Glimell (2011) for a synthetic historical panorama.

innovative ideas (e.g Graves et al. 2011; Haufe 2013; Boudreau 2016; Fang and Casadevall 2016; Vaesen and Katzav 2017²). The recent proliferation of proposals to rethink the concrete conditions of science funding is strongly motivated by natural scientists themselves, and by the literature on the economics of science. But how may philosophy of science and epistemology enter these debates? Our paper considers this question by proposing a critical assessment of the benefits and limits of the social epistemology models of “division of cognitive labor”, when trying to deliver policy advice.

One of the main epistemological questions at stake about the mechanisms of science funding concerns the consequences of the institutional arrangements which regulate the allocation of resources to scientific research on the growth of knowledge. Formulated as such, our first interrogation exhibits its link with what Goldman and Blanchard (2016) defines as the “third branch” of social epistemology³. As noted by the authors, this “third branch” largely focuses, since Kitcher’s (1990) early work, on the question of the *division of cognitive labor* (DCL), seeking to describe what might be an optimal division of tasks between competitive approaches within a given scientific community, and to assess the institutional conditions promoting the achievement of this optimum. More precisely, Kitcher’s (1990) initial question is the following: “what is the optimal division of cognitive labor within a scientific field, and in what ways do personal epistemic and non-epistemic interests lead us toward or away from it?” (p.22). This inquiry was recently renewed by Viola (2015) which asks “(1) which is the optimal distribution of cognitive efforts among rival methods within a scientific community? and (2) whether and how can a community achieve such an optimal distribution?” (p.1). These questions ground an active line of research, mostly dedicated to the development of mathematical models. Such models propose solutions to Kitcher’s optimization problem by focusing on rewards or credit schemes (e.g Strevens 2003; Zollman 2018), on the cognitive structure of the scientific community (e.g Weisberg and Muldoon 2009), or on the effect of project selection by centralized funding mechanisms (Avin 2018). Since recently, some of

²Of course, the question of the governance of science is not limited to that of the processes of grant allocation by funding agencies. Our paper exclusively deals with epistemic considerations, but ethics and political philosophy also have their say. The so-called “responsibility” of research and innovation (Arip 2016), and the democratic involvement of citizens in science policy decisions (Kitcher 2001) constitute central perspectives on this matter.

³Devoted to the “assessment of the epistemic consequences of adopting certain institutional arrangements or systemic relations as opposed to alternatives”.

these agents-based models manifest an explicit interest in providing concrete policy advice (Viola 2018), notably about the modalities of the allocation of resources (Kummerfeld and Zollman 2016; Avin 2018) and the cognitive structure of scientific communities (Pöyhönen 2016). However, this ambition is tempered by the existence of inherent limits to modeling activities in social epistemology. As with models designed to describe a domain of the real world, it is possible (and welcome) to question the simplifications, idealizations and background hypotheses which are susceptible to amend the descriptive and normative value of DCL models (Yloski 2014; Pöyhönen 2016). Given the well-recognized difficulties to assess the exact benefits and limits of these works, it is then common to note that “there is much work left to do” (Muldoon 2013, p. 124) to improve or judge their epistemological and political value. The injunctions to prudence in interpreting the results of the DCL models⁴ raises the question of their exact interest in guiding the institutional regulation of research, and notably the mechanisms of grant allocations by funding agencies.

In this paper, we aim at showing that a major threat to the descriptive and normative value of current DCL models is a fundamental lack of clarity about the exact *object* which is divided. As we shall see, authors consider without distinctions the allocation of resources to “theories”, “approaches”, “methods”, “research programs”, “research projects” etc. By doing so, they bypass two important dimensions of the epistemological reflection about the division of cognitive labor. First, the epistemic choices made by researchers are determined by a hierarchical network of shared objectives, from experimental, technical practices to conceptual and representational activities. Second, the significance and the form taken by epistemic diversity are highly dependent on the *kind* of object under study. Our argument is that these elements should be carefully taken into account when building DCL models and when interpreting the results obtained. We suggest, based on concrete instances of such social epistemology modeling, that the ignorance of these features leads to unsatisfying conclusions about the desirable institutional conditions of research. Finally, more than a refinement of the

⁴Some instances of this interpretative prudence, among others: “Can our models fix useful concepts and provide templates for causal mechanisms that could be at play? Could they be used to help shape the debate around emerging policy decisions? The answers will come from future work in the field” (Avin (forth.), p.32); “The cautious conclusion to be drawn from these differences is that, in its entirety, the relationship between diversity and epistemic performance is likely to be more complex than can be captured by any simple model” (Pöyhönen 2016, p. 4530); “Of course these models are limited in two critical ways. First, they proceed into idealizations about the structure of the scientific community and about individual scientists. Real scientists and scientific communities are more complicated than our models, and it is always possible that a critical causal factor has been left out” (Zollman 2018, p. 26).

mathematical models used, what is needed to improve social epistemology insight in research policy is a *qualitative* reflection about the dynamics of scientific progress.

2-DCL models: common architecture

a-A (short) systematic review

Many reviews of the different kinds of works dealing with the DCL problems are available (see, for instance, Muldoon 2013; Goldman and Blanchard 2016; Avin (forth.)). Let us reiterate here the main directions taken by philosophers interested in this question since Kitcher's (1990) seminal study. The starting point is the idea of a tension, within scientific communities, between individual and collective rationality. Kitcher imagines the case of a shared objective (for instance, determining the physical structure of a certain molecule), with two competitive methods to reach it – method I and method II, method I being known as more accurate. For Kitcher, a purely rational epistemic agent will choose method I to solve the problem posed. Yet, this behavior does not optimize the division of cognitive labor, which would (mathematically) benefit from a more balanced distribution of human resources between method I and method II. Kitcher then shows that this idealized community may be nearer to the optimum if there exists a reward scheme according a larger reward to scientists working on the less popular “research program”. This mechanism will be efficient on the condition that scientists do not follow a strictly epistemic rationality, but have other sources of motivation (professional credit or rewards). Following this line, Strevens (2003, 2013) proposes another reward scheme, based on the “priority rule” according to which the first research program that discovers a certain result gets all the reward. Zollman (2018) develops a quite distinct mathematical model to assess the effect of the search for professional credit (*versus* the “seek the truth”) on scientific progress.

Besides rewards and credit, other factors were studied for their ability to optimize the division of cognitive labor. Zollman (2010) addresses the formation of consensus on the right theory of a given phenomenon. Considering two theories T1 and T2, T1 being favored in the initial experiments, he suggests on the basis of historical instances that it may be rational, at the scale of a given scientific community, to pursue the poorly justified theory T2 (because T2 may be, finally, the right one and should not be eliminated too quickly). Using computer simulations, he then (not surprisingly) shows that the communication rate between individuals

influences the collective rationality: structures with *less* communication may score better in choosing the right theory.

A third kind of factor influencing the division of cognitive labor is mobilized by the epistemic landscape models launched by Weisberg and Muldoon (2009). Here, it is the “cognitive diversity” which is put to the fore. The authors consider an “epistemic landscape” explored by individuals following distinct engagement strategies. This epistemic landscape represents, in a three dimensional space, all the “research approaches” possible in a given field of research (Pöyhönen 2016), and their respective significance⁵. Epistemic agents may adopt a “control” engagement strategy (they move in the landscape without taking into account the behavior of the other agents), or be “followers” (they adopt the most significant programs explored by their predecessors); they also can behave as “Maverick”, and explore unknown zones. The idea is then to determine, depending on the initial conditions (size and shapes of the landscape) which population, or mix of populations, optimize the division of cognitive labor (the cumulated significance of the research process) (Weisberg and Muldoon 2009; McKenzie and Himmelreich 2015; Pöyhönen 2016).

All these works share an identical macro-conception of the dynamics of science as a closed system internally regulated by an invisible-hands mechanism (Wray 2000). Individual scientists are free to choose the direction of their research, in a context of unlimited resources (Viola 2015). In particular, there is “no superintendents” (Strevens 2013, p. 21). Against this idealization, Viola (2015) insists on the necessity “to consider the role of external social factors such as the political decisions to pursue some or other scientific project” (p. 9). In particular, the role of funding agencies should be explicitly addressed. Following this line, Avin (forth.) renews the epistemic landscape model to study the effect of “centralized funding” on the division of cognitive labor. He simulates distinct ways of funding science, based on the estimated significance of the project (its position on the epistemic landscape), on the time passed by the individuals in the system, or on a lottery (random distribution). Depending on the size of the epistemic landscape, he shows that the lottery strategy may be the best one in terms of cumulative significance of the projects followed. He concludes that dealing with the classical opposition between “plausibility and originality” (Polanyi 1962) or

⁵As noticed by Pöyhönen (2016), significance may be understood according to Kitcher (1993) (Chap. 4) and Kitcher (2001), as “the significance of the truths that can be uncovered by using this approach” (Pöyhönen (2016), p.4522).

“exploration and exploitation”, a random allocation of resources may optimize the division of cognitive labor. He links this result to more qualitative arguments defending lottery as a good way to select projects (Gillies 2014, Fang and Casadevall 2016).

Let us add, to close this rapid outlook on the DCL models, the work of Kummerfeld and Zollman (2016), which aims to quantify the “conservatism” of a scientific community constituted by individuals left to “their own devices” (p. 1057). As Kitcher did, but by using more intricate mathematical machinery, they suggest that individual epistemic rationality conflates with the optimum division of cognitive labor when individuals have to choose between a “risky” and a “safe” alternative. They conclude that the mechanism of grant allocation should voluntarily fund a certain amount of “risky” projects, to compensate the endogenous “conservatism” of the scientific community.

All these DCL models have a common architecture, that is to say, they share a common general formulation of the question of the division of cognitive labor in science, and a common scheme to solve it. The starting point is to consider a certain pre-defined objective O to reach, which is shared by a given community. Typically, this objective may be a (theoretical, experimental, utilitarian, technical) problem to solve. It may also be, specifically in the case of epistemic landscape models, an ensemble of objectives defining a research area. The authors then consider the existence of various means M_i to reach O , each of them having a *utility function* quantifying their ability to fulfill O . The central aim is then to distinguish the individual instrumental rationality⁶, which tends to maximize the probability, for each researcher, to reach O , and other “non-epistemic⁷” factors, such as the search for rewards or

⁶Let us maintain, following Kelly (2003), that instrumental rationality classically designates “the rationality which one displays in taking the means to one's ends” (p. 612)”. It is opposed to *epistemic rationality*, defined as “the kind of rationality which ones displays when one believes propositions that are strongly supported by one's evidence and refrains from believing propositions that are improbable given one's evidence”. In DCL models, as noted by Zollman (2018), epistemic rationality is understood as a form of instrumental rationality: “epistemic rationality is a species of instrumental rationality, viz. instrumental rationality in the service of one's cognitive or epistemic goals” (Kelly (2003), p. 613). In other words, epistemic rationality is expected to be mobilized to choose the *best* alternative (the M_i which has the estimated highest probability to reach O) to solve a cognitive or epistemic problem. This instrumentalist conception of epistemic rationality corresponds to the multiplicity of the (cognitive, practical, utilitarian) kinds of objectives considered in the DCL models.

⁷The opposition between “epistemic” and “non-epistemic” factors may be discussed, since the objective of DCL models is precisely to show that “non-epistemic” factors positively influence the search for truth. However, the rationale behind this distinction lies on a contrast between the *motivations* guiding the choices made by researchers, as well exposed in Zollman (2018). If these motivations are *exclusively* those of solving the problems posed to the community, they are considered as acting as purely epistemic agents –even if this problem is not itself strictly cognitive.

professional credit, the cognitive features of the individuals or the centralized selection of means by funding agencies. The shared general conclusion, since Kitcher's work, is that researchers (or peer-reviewers affiliated to funding bodies) only following individual instrumental rationality do not optimize the division of cognitive labor. Indeed, since all possible research projects (M_i) have a certain probability to reach O , it is counter-productive to concentrate resources only on the most promising approaches. Yet, individual instrumental rationality is supposed to generate the phenomenon of herding on secure alternatives. Reciprocally, "non-epistemic" factors may positively modify the division of cognitive labor by promoting the exploration of more risky alternatives.

The first step in the elaboration of DCL models is then to define a *theoretical* optimal distribution of resources. This optimal DCL may be mathematically expressed (e.g in Kitcher 1990; Strevens 2003; Kummerfeld and Zollman 2016), or implicitly fixed by the initial conditions chosen. This is the case for the epistemic landscape models, where the pre-determined form of the landscape determines what would be an optimal distribution of labor to explore it given certain initial funding conditions (number of researchers and number of projects selected) and respect to a given variable (for instance the cumulative epistemic significance reached).

b-What exactly is divided?

Let us specify this general scheme, by addressing a crucial point: what, exactly, is thought to be divided? In other words, what *are* exactly M_i and O ? The question is relevant, since the authors apply the same schemes to distinct objects. As noted by De Langhe (2014), "Kitcher's basic unit of analysis is "methods," Strevens uses "research programs," and Brock and Durlauf use "theories."". He adds: "because as far as their dynamics of adoption are concerned the literature *uses these concepts interchangeably*, I will use only the concept of "theories" understood as standards for the division of cognitive labor to which individual scientists make contributions" (p. 445, *our italics*). Kummerfeld and Zollman (2016) explicitly recognize this indeterminacy, when they clarify their notion of "general research project" by including in it "different theoretical commitments, paradigms, research methodologies, treatments strategies in medicine, and so on" (p. 1059). This vagueness of the object which is divided is particularly significant in the epistemic landscape models. Pöyhönen (2016) explicitly addresses the question of "what does an epistemic landscape represent", and deduces that they are *not* "a search space for a single problem" (p. 4525). An

epistemic landscape would stand for a “scientific research topic (e.g synthetic biology, astrophysics, endocrinology)”, divided into “different but complementary research approaches”. For instance, the authors argues, “attempting to synthesize novel DNA nucleotides and studying the stability of these molecules by computational methods are independent but both necessary research approaches in synthetic biology” (p. 4523). Interestingly, he adds that the “discrete patches” composing the epistemic landscape “represent a combination of (i) a research question being investigated, (ii) instruments and methods for gathering and analyzing data, and (iii) background theories used to interpret the data”. Consequently, as the author confesses himself, “epistemic landscapes underlying real scientific research involve a greater number of interdependencies between the elements of approaches (question, instrument, methods, theories)” (p. 6). These precisions are the most complete we may find within the literature of epistemic landscape models. The works of Weisberg and Muldoon (2009) or Avin (forth.) simply consider epistemic landscapes as an ensemble of “research approaches” or “projects”.

This diversity of *Mi*'s nature also reflects in the multiplicity of kinds of objectives *O* which are mobilized as examples illustrating the DCL models. Kitcher (1990) takes the case of the elucidation of the physical structure of a given molecule (of medical importance); Zollman (2010) that of the explanation of a disease (the peptic ulcers), or more generally that of all “truth seeking” objectives (Zollman 2018); Goldman and Blanchard (2016) cites the elucidation of the structure of the DNA; Pöyhönen (2016) considers “the study of opioid receptors in chemical biology, or critical phenomena in statistical physics” (p. 4530) etc.

What we want to show in the next section is that because of the interchangeable use of distinct kinds of ends and means, two important dimensions of the dynamics of scientific progress are *absent* from the DCL models. Then, in the third section, I will suggest that these *absent* dimensions are *missing* dimensions, in the sense that they should be explicitly taken into account when interpreting the results of DCL models –which is never done in the corresponding papers. We conclude that to be politically relevant, works on DCL do *not* need to improve the sophistication and/or the mathematical complexity of their models, but should consider seriously some qualitative epistemological issues about the logic of scientific research.

3-Two absent dimensions of the debate

The argument we will defend here is that the current DCL models may not give an accurate description of what could be a theoretical optimal division of cognitive labor (first step of the DCL models). In particular, two dimensions are absent. First, it has to be noticed that the scientific field is structured as a system of hierarchical, interconnected practices motivated by a superposition of different kinds of objectives. Second, the epistemic significance of pluralism is highly dependent on the *nature* of the object under study, and as a consequence, depends on the discipline or field of research considered.

a-Research objectives are embedded within various systems of practices

Let us imagine a given objective O that could be taken as an example in a possible DCL model: the search for a treatment against a new, unknown bacterial disease D . Let us also suppose that there exist two general ways to solve this problem: for instance, the search for a vaccine (Mv), or for a new antibiotic molecule (Ma). Each mean Mi has an estimated utility function, and epistemic agents are supposed to distribute their effort following different rules (instrumental rationality, credit, rewards, random selection etc.). These utility functions may be estimated, for instance, on the basis of past successes in dealing with bacteria belonging to the same biological class. However, this is not the end of the story. We can imagine that each Mi constitutes itself an objective which can be fulfilled by following distinct possible ways. For instance, there may exist various different methods to prepare a vaccine (by targeting distinct proteins from the bacterial membrane), and various techniques to find antibiotics (molecular screening, or rational-drug design (Adam 2005)). And each of these methods defines a new objective, with potentially various means to achieve it. This recursive logic may be even more complex if we consider higher-level alternatives. For instance, the modern techniques and concepts used in medicine, anchored in cell and molecular biology, may be questioned by more *holistic* conceptual frames –found in some “traditional” medical practices. The initial objective O may then be fulfilled within two general approaches, anchored in opposite representations of diseases, organisms, or life itself ($C1$ and $C2$). Besides, it is worth noting that the competition between $C1$ and $C2$ may have an importance for many other objectives O –in medicine or fundamental biology. To consider this interlocking of objectives, we propose to invest, as a conceptual toolbox, the general “grammar of scientific practices” recently formulated by Chang (2014).

The contemporary attention to science “in practice” (Soler et al. 2014) gives a representation of science, as a process, made of interconnected networks of more or less independent types of practices. The so-called “new-experimentalism”, notably launched by Hacking (1983), aims at describing in detail the heterogeneous elements entering experimental practices (Chang 2014). These experimental practices are considered as “hav[ing] a life of their own...” (Hacking 1983, p. x), independently of their hypothesis-testing role. Later, the need appeared to integrate *representational* (conceptual, theoretical) activities in the frame of scientific practices (Woody 2014). Following this line, the activities of epistemic agents are multidimensional, since this notion of practice includes “physical, mental, and “paper-and-pencil operations””, as argued by Chang (2014) (p. 68). Chang then proposes the notion of “system of practices” to describe the hierarchical interlocking of “epistemic activities” motivated by distinct “aims”. First, in this scheme, “all scientific work, including pure theorizing, consists of *actions*” (*our italics*, p. 67). An action (or an activity, in Chang’s “grammar of scientific practices”) performed by an epistemic agent, is characterized by the “presence of an identifiable aim” (p. 72). To reach this aim, the agent develops “a coherent set of mental or physical operations (...) in accordance with some discernible rules” (*Idem*). This activity, led by “inherent purposes”, is also motivated by “external functions”. For instance, the act of lighting a match mobilizes a set of operations, and may be part of a more global plan (for instance, to light a Bunsen burner), itself constituting an activity within a larger system of practices (aiming to answer a theoretical question, or to find a new therapy etc.). The notions of “activities” and “systems of practices” are then relative, and depend on the scale of analyses. The important point is that each action performed by an epistemic agent is integrated in a hierarchical network of purposes, and is characterized by its internal *coherence*⁸. In Chang’s scheme, coherence is ensured by the *success* in achieving one’s end. More generally, we may consider that the coherence is linked to the relevance of the means used to reach the objective motivating the activity (or system of practices) in question. In other words, coherence is ensured by the deployment of instrumental rationality, leading to determinations of the means *most* adapted to the ends considered (Kelly 2003; Zollman 2018). Let us recall here that these activities or systems of practices may gather distinct *kinds* of (mental, physical) objectives, and then distinct kinds of means *Mi*: technical operations, setting-up of experimental protocols, formulation of new concepts or new theories, etc.

⁸For a graphical representation of this hierarchical imbrication of activities, see the “commentary” on Chang’s work by L. Soler and R. Catinaud in Soler et al. (2014).

How can we apply this scheme to the question of the division of cognitive labor? As DCL models do, let us first consider, in a God's eye perspective, the possibility to define a theoretical optimal division of cognitive labor given a certain problem to be solved (objective O). This resolution may be achieved by developing distinct activities. These activities are integrated in a system of practices, and are then defined both by their *inherent* purposes and by their various hierarchical *external* functions. Given this situation, two remarks should be made. First, inherent purposes may be divided into a variable number of underlying objectives. For instance, if we refer to the instance proposed at the beginning of this section, the search for a treatment to disease D may be considered as the general end of a system of practices including a first level of sub-objectives o and o' (the search for a vaccine or for an antibiotic treatment). These constitute the inherent purpose of sub-systems of practices (including technical, experimental, or theoretical operations). We can expect that depending on the nature of O (technical, experimental, theoretical etc.), the thickness (the number of imbricated systems of practices) of the corresponding network may be highly variable. Finally, each DCL optimization problem is best described as a *superposition* of optimization problems.

Second, a given activity may have various external functions. For instance, the development of an experimental technique T (e.g a method of protein extraction from bacteria) may be important for various possible alternatives $Mv1$ and $Mv2$ in preparing a vaccine against D . Besides, this experimental technique may be relevant for other systems of practices –in immunology, biochemistry, clinical medicine, or molecular biology. Let us imagine now that $Mv1$ and $Mv2$ (the “vaccine” solutions to cure D) have a lower utility function than Ma precisely because the protein extraction activity is not efficient. If we consider only $Mv1$, $Mv2$ and Ma , the optimization problem will have a certain theoretical solution. This solution would be very different if we posed the optimization problem between Ma and T : we can expect that T would have a higher utility function than both $Mv1$ and $Mv2$. This is true because T is supposed to be important for $Mv1$, $Mv2$, and for other systems of practices. This last formulation would be more accurate with respect to the real state of the system of scientific practices. Finally, the solution of a given optimization problem varies depending on the scale of the analysis. “Theories”, “methods”, “approaches”, “programs” are not interchangeable units. Ideally, a well-posed DCL problem should consider *altogether* these distinct

dimensions, reflecting the hierarchical interlocking of objectives characterizing the systems of scientific practices.

b-The management of plurality may not be always considered as an optimization problem

In the multiplicity of kinds of objectives and means taken interchangeably as basic units in DCL models, distinct disciplines are cited: physics (the study of critical phenomena in Pöyhönen 2016), chemistry (the elucidation of the structure of a given molecule or the opposition between phlogiston theory and modern chemistry in Kitcher, 1990), fundamental biology (the structure of DNA in Goldman and Blanchard, 2016), medicine (the explication of peptic ulcers in Zollman, 2010). These instances designate different kinds of objectives (theoretical, technical, experimental), but they also concern various kinds of objects. Yet, we argue that this generalization is biased since many scientific objects generate questions whose resolution cannot be considered as a common DCL problem. To make this point, we propose to consider the notion of *complexity*, classically defined as the co-existence of multiple causal pathways, belonging to distinct levels of organizations (Mitchell 2009). For instance, mental pathologies, or human behavior more generally, constitute typical complex phenomena, with genetic, biological and environmental causes possibility non reducible to the molecular level. In these cases, the question of the multiplicity of radically distinct *approaches* might not be formulated into an optimization problem: the different directions of research are better described as distinct heuristics than as alternative means characterized by their estimated utility functions.

Let us take, for instance, as an objective *O*, the seek for an explanatory theory of a particular behavioral (normal or pathological) trait. Following Longino (2013), we can consider the co-existence of distinct kinds of disciplines that she identifies as “quantitative behavioral genetics”, “socio-environmental approaches”, “molecular behavioral genetics”, and “neurobiological approaches” (p. vii). At the metaphysical level, this co-existence may either be considered as a (temporary or permanent) consequence of our cognitive limitations, or as a result of the existing diversity of levels of organization in nature (Ruphy 2005). In all cases, it is arguable that this plurality has an inherent value which exceeds the estimated efficiency of each approach to provide a good explanation to the behavioral trait we consider. This is due to the fact that, at least temporarily, all the available explanations may capture *something* of the phenomenon under study (for instance, a partial causal mechanism). In this frame, even a quite marginal alternative (for instance, psychoanalysis), may deserve to be pursued, because

it may constitute a possible *heuristic* allowing the identification of particular causal pathways. This idea was defended in the case of theories of cancer. The classical, gene-centered theory, is currently challenged by an “organicist” one (the Tissue Organization Field Theory, Soto 2011), and these competitive approaches may be considered as two distinct heuristics to exhibit the causes of cancer (Malaterre 2007). Finally, we argue that for complex objects, the co-existence of distinct approaches cannot be thought in the frame proposed by current DCL models. The alternatives may not be compared in their efficiency to reach a pre-defined objective. Indeed, if we accept that pluralism reflects, to some extent, the variety of causes determining the phenomenon at stake, then these alternatives are mutually dependent (each approach delivering *a part of* the explanation). The situation is distinct in most of the instances provided in the DCL literature, where each approach is supposed to have a singular, independent utility function.

In the next section, we aim to show that these two *absent* dimensions (the hierarchical interlocking of objectives and the ontological limits to the DCL formalism) are *missing* dimensions. They would need to be explicitly addressed in order to interpret in a correct way the results provided by the DCL models.

4-Why these absent dimensions are missing dimensions

As most of the authors writing on DCL models, we are conscious that the very modeling activity inevitably presupposes the use of idealizations or simplifications which do not preclude *in themselves* their heuristic value –their ability to capture *certain* features of a given phenomenon (Muldoon and Weisberg 2011; Ylikoski and Aydinonat 2014; Pöyhönen 2016). However, this heuristic value tightly depends on a careful work of interpretation of the results they provide. Here, we argue that the negligence of the elements of the dynamics of scientific research we considered in the previous section may generate important misinterpretations when DCL models try to provide advice for science policy. We think the best way to defend this thesis is to consider some concrete examples of DCL models whose conclusions aim to formulate political insights. By doing so, our objective is to show why these absent dimensions may constitute, very concretely, missing dimensions.

a-Example 1: funding science by lottery?

The first instance we will consider is that of Avin's (forth., 2018) case in favor of the introduction of some random elements in the mechanism of grant allocation to scientists. Apart from its synthesis of qualitative arguments defending random allocation as a possible way to overcome the "conservatism" of peer-review (Gillies 2014), Avin proposes a contribution through a DCL model. Based on this model, he suggests that, given a certain landscape progressively explored by epistemic agents, a peer-review only based on instrumental rationality (selecting projects on the basis of their estimated significance) might be, in some cases, sub-optimal. On the contrary, a system of lottery which randomly allocates resources may out-perform the classical peer-review process. Intuitively, this result is directly linked to the fact that the epistemic landscape is not known in advance: all estimation of the future significance of a project may then be mistaken. On the basis of this result, he proposes a mechanism where a certain proportion of projects (those which do not belong to the x% best or worst proposals as evaluated by a first round of selecting peer-review) enter a process of random funding. We argue that this proposal is biased, or insufficient, since it is grounded on a misrepresentation of the very structure of the systems of scientific practices. The interest of random allocation is assessed, in Avin's model, for the cases of poorly known epistemic landscape, with a great number of projects whose potential significance is difficult to estimate. In these (hypothetical) situations, it is quite intuitive indeed that a random distribution may be the most efficient way to *explore* the space. Avin (2018) links these situations to "basic research" (p. 33). In parallel, he lists some situations where "lottery should not be used" (p. 31). Among them, he evokes the case of projects with "bounded uncertainties" (where the need to obtain a quick answer, for instance under the pressure of "external constraints", prevents "any significant exploration of uncertainties or open-ended avenues") and of "fully explored" epistemic landscapes (where the knowledge of a given "area" is sufficient in yielding a good estimation of the utility functions of the project to be pursued). It seems to us that these two cases express the same general (qualitative) idea: when the problems to be solved are *sufficiently* clearly defined, and the significance of the possible research project to address them are *sufficiently* known (to sum up: when the epistemic landscape is sufficiently known), then random allocation might *not* be the optimal solution to divide cognitive labor. Following this criterion, we argue that it is not at all obvious that there do really exist "epistemic landscapes" adapted to such a funding by lottery. Indeed, the rationale behind Avin's research policy advice lies on the assimilation of "basic research" to large epistemic landscapes with many "peaks" whose significance is strongly indeterminate. In Chang's terms, these peaks represent possible activities, with both inherent purpose and

external functions. As we noticed in the previous sections, these activities are embedded in systems of hierarchical, interconnected practices. The *global* significance of a given activity is then dependent on the *partial* significance of all the practices constitutive of its inherent purpose, and of the others systems of practices which would benefit from its development. Let us imagine a project *P* with a poorly estimated significance in a given “epistemic landscape” *E*. *P*, and the sub-activities it entails (e.g. the development of new techniques or experimental procedures), may have a clear significance for other systems of practices, corresponding to other “epistemic landscapes”. Its sub-activities may also have known significance within *E* itself, independently of *P* (for instance, if they constitute sub-objectives of other projects whose significance is better known). It is worth noting that *P* or its sub-activities may also be important to face some pressing external constraints – situation where, following Avin himself, lottery is not welcome. Finally, we argue that lottery’s argument is based on the *hypothetical* existence of many *isolated* projects (activities or systems of practices), whose significance would be strongly indeterminate. We think that the qualitative assessment of the structure of the systems of scientific practices suggests that the very existence of such “epistemic landscapes” needing a funding by lottery is not at all self-evident. Consequently, we argue that the recognition of the network structure of scientific activities leads one to formulate research policy advice opposed to Avin’s. Random allocation is presented as a solution to the lack of information about the significance of proposed research projects. Our alternative view suggests that a more exhaustive examination of the relevance of each research project with respect to the existing systems of practice would show that there are fewer cases than Avin anticipates in which the epistemic landscape is insufficiently known. This examination would also show that when the epistemic landscape is insufficiently known, the degree to which it is unknown is overestimated by Avin. Thus, the need for randomized funding allocation is mitigated, and perhaps evaporates entirely⁹.

The important point is that even if Avin’s model is instructive and well-suited to describe the way science is funded, its normative aim (guiding the way science *should* be funded) might be impaired by the conception of the essential properties of the research process it is based on. Our argument may open another direction of investigation into the optimization of grant allocation through peer-review processes. As we suggested, one of the main epistemic challenges of science policy is the evaluation of the comparative interest of the proposed

⁹This exhaustive examination may be completed, in the spirit of Kitcher (2001), by a democratic assessment of the various desires expressed by ordinary citizens.

projects. The relevance of this evaluation depends, indeed, on the knowledge we have of the insertion of these projects into the existing hierarchical network of interconnected practices. Yet, as rightly pointed out by DCL models, and notably by Avin's model, the peer-review process may not be well-suited to adequately capture the current relationships between these systems of practice. The empirically measured lack of robustness of peer-review confirms this point (Graves, 2011). However, rather than through a random distribution of resources, we argue that a better evaluation of the interest of the proposed projects could be reached through a more decentralized mechanism of evaluation. As a limited panel of reviewers does not represent well the current state of scientific practices, an active participation of all the scientists actively practicing research in choosing the projects to be funded could provide a better knowledge of the objective interest of these projects. For instance, we propose that each scientist could have a right (and a duty) to choose a limited number of projects among the ensemble of all the propositions made by their colleagues; the projects presenting the best average mark would then be funded¹⁰. We argue that with this kind of global scheme, the convergence of interests (and so the estimated value of a given project) would be better evaluated. Obviously, this decentralization of evaluation is susceptible to be fully efficient only if there exist classificatory systems enabling each researcher to easily identify the projects which are interesting to him. Without giving a complete practical solution, we suggest that the use of key-words or key-expressions, similar to that developed to classify scientific literature in numerical bases of data, may be a fruitful direction to follow.

This proposal may open interesting modeling possibilities. In all cases, DCL models, if they want to improve their utility in designing science policy strategies, should represent the real multidimensional structure of the systems of scientific practices in a more suitable way.

b-Example 2: promoting "risky" strategies?

A second instance of DCL models proposing explicit political insights may be found in Kummerfeld and Zollman (2016). The core of their model lies on a distinction between "safe" and "risky" alternatives, which may be "different theoretical commitments, paradigms, research methodologies, treatment strategies in medicine, and so on" (p. 1059). These

¹⁰ We may also imagine here that publics which are exterior to the scientific field (citizens, economic sphere or political actors) could also have their say in this process of proposing and voting for projects.

alternatives have utility functions with respect to given, pre-determined, objectives¹¹. In Kummerfeld and Zollman's model, a "safe" alternative is known to be the "best line of research", that is to say the one "that has given the highest average payoff so far" (p. 1062). The risky one is considered as such because its (Gaussian) utility function, even if it has a higher mean value, is poorly known. The simulation machinery shows that when the utility function of the "risky" alternative has a higher mean value than the utility function of the safe one, the individual instrumental rationality is sub-optimal. By neglecting the risky alternative, individuals do not maximize the global utility of the community when distributing their cognitive effort.

The authors deduce from this result that funding agencies should actively "encourag[e], in *some situations*, unpopular, risky science" (p. 1057, *our italics*). The central challenge to make this conclusion politically relevant is to delineate more precisely these "situations" in which uncertain projects should be voluntarily funded. We argue that this task needs to take into account the hierarchical interlocking of objectives structuring the systems of scientific practices.

Let us note that the authors consider cases where a given choice is made to fulfill one, and only one, well-defined objective. Yet, as we have discussed previously, each "research project" at stake may have multiple external functions (Chang 2014). This is equally true for the sub-objectives implied by each alternative. To revisit our previous example, the development of a new vaccine to cure the disease D (a typical "research project" in Kummerfeld and Zollman sense) may be useful in other areas. For instance, it may lead to the chemical and physical characterization of a given protein P of the bacterial membrane, which is important in immunology, cell biology etc. It may also lead to the improvement of a given technique T to extract proteins from bacterial cells. We can suppose in this case that depending on the objective we consider, the "research projects" will have *distinct utility functions*. Even if the development of a vaccine is a "safe" alternative with a relatively low utility function (in comparison to, say, the search for antibiotics) to cure D , the knowledge of protein P and the improvement of T may have high utility functions relatively to other systems of practices (motivated by distinct general objectives). In this case, the herding of scientists in the "safe" alternative may be counter-productive for the problem at stake (the

¹¹If we follow the examples provided by the authors, these objectives may be to cure peptic ulcer, to make astronomical predictions of the "locations of heavenly bodies" in 1550 (by choosing between Ptolemaic or Copernican paradigms, p. 1063) etc.

search for a treatment against D), but positive at a larger scale (relatively to other systems of practices). Kummerfeld and Zollman's argument, if valuable in local contexts (that is to say, to solve a precise problem, or to explore a given, well-defined epistemic landscape), is harder to justify in more global ones (when considering the existing interconnections between the systems of practices). Finally, the apprehension of the situations where "unpopular" science should be actively promoted is certainly far from being an easy task.

It could be noticed against our argument that Kummerfeld and Zollman's model, such as the majority of DCL models, aims to study the optimal distribution of cognitive effort precisely in such local contexts, where each scientific objective is thought in isolation from the others. However, from the point of view of research policy, the optimization of the division of cognitive labor is a global problem, where all disciplines, specializations, research questions are to be thought together. Consequently, in their ambition to provide political insights, DCL models should mobilize a more realistic account of the structure of the systems of scientific practices.

5-Conclusion

The question of the relevance of social epistemology to guide research policy is a pressing one, as showed by the recent publishing of a volume of the *Roar Transactions* on this subject (Viola 2018). In this paper, we address it in the specific case of DCL models, which have flourished since the publication of Kitcher's (1990) seminal work. These works try to assess what exactly is an optimal distribution of cognitive labor among researchers, and how the institutional conditions regulating science may contribute to its achievement. We show that when faced with the first question, DCL models ignore two central dimensions of the dynamics of science (the hierarchical interlocking of objectives and the variable epistemic significations of pluralism). We argue that this ignorance is problematic, since it may lead to misinterpretation of the results of the DCL models, and to biased policy advice. In this sense, they really constitute missing dimensions in the debate. This conclusion does not lead to a rejection of DCL models *per se* as epistemologically and politically irrelevant, but aims at figuring out some important properties of the dynamics of science that should be taken into account in this modeling activity.

As Muldoon (2013) confesses, “there is much work left to do” (p. 124) to understand the “benefits and burdens of diversity” and the optimal division of cognitive labor. What form should this work take? We argue that the search for more refined quantitative DCL models should be supplemented by qualitative studies of the epistemic or psychological principles guiding the social organization of science, the consequences of the pressure exerted on individual scientists by centralized piloting, the balance between targeted and free research are crucial issues to inform science policy, and qualitative history and philosophy of science certainly have a central role to play on this matter.

References

Adam, M. 2005. Integrating research and development: the emergence of rational drug design in the pharmaceutical industry. *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 513-537.

Arip, A. 2016. The clothes of the emperor. An essay on RRI in and around Brussels. *The Journal of responsible innovation* 3(3): 290-304.

Avin, S. Centralised funding and epistemic exploration. Forthcoming in *The British Journal for the Philosophy of Science*.

Avin, S. 2018. Policy Considerations for Random Allocations of Research funds. *Roar Transactions* 6(1).

Boudreau, K. J., Guinan, E. C., Lakhani, K. R. and Riedl, C. 2016. Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science. *Management Science* 62(10): 2765–2783

Chang, H. 2014. Epistemic Activities and Systems of Practice: Units of Analysis. In *Philosophy of Science After the Practice Turn*, ed. Soler, L., Zwart, S., Lynch, M., and Israel-Jost, V. Routledge.

De Langhe, R. 2014. A unified model of the division of cognitive labor. *Philosophy of Science* 81: 444-459.

Fang, F-C. and Casadevall, A. 2016. Research Funding: the Case for a Modified Lottery. *mBio* 7(2): e00422-16

Gillies, D. 2014. Selecting applications for funding : why random choice is better than peer review. *RT. A Journal on research policy and evaluation* 2(1).

Goldman, A. and Blanchard, T. 2016. Social Epistemology. In *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta, winter 2016 edition.

- Graves, N., Barnett, A. G. and Clarke, P. 2011. Funding grant proposals for scientific research: retrospective analysis of scores by members of grant review panel. *BMJ* 343.
- Haufe, C. 2013. Why do funding agencies favor hypothesis testing? *Studies in History and Philosophy of Science* 44: 363-374.
- Kelly, T. 2003. Epistemic Rationality As Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research*. LXVI(3): 612-40.
- Kitcher, P. 1990. The division of cognitive labor. *The Journal of Philosophy* 87(1): 5-22.
- Kitcher, P. 1993. *The Advancement of Science*. Oxford University Press.
- Kitcher, P. 2001. *Science, Truth and Democracy*. Oxford University Press.
- Kummerfeld, E. and Zollman, K-J-S. 2016. Conservatism and the scientific state of nature. *The British Journal for the Philosophy of Science* 67(4): 1057-1076.
- Longino, H. 2013. *Studying Human Behavior. How Scientists Investigate Agression and Sexuality*. The Unversity of Chicago Press.
- Malaterre, C. 2007. Organicism and reductionism in cancer research : Towards a systemic approach. *International Studies in the Philosophy of Science* 21(1): 57-73.
- Mitchell, S. 2009. *Unsimple truths: science, complexity, and policy*. The Unversity of Chicago Press.
- McKenzie, J. and Himmelreich, A.-J. 2015. Epistemic landscapes, optimal search, and the division of cognitive labor. *Philosophy of Science* 82: 424-453.
- Muldoon, R. 2013. Diversity and the division of cognitive labor. *Philosophy Compass* 8(2): 117-125.
- Muldoon, R. and Weisberg, M. 2011. Robustness and idealization in models of cognitive labor. *Synthese* 183(2): 161-174.
- Polanyi, M. 1962. *The Republic of Science: Its Political and Economic Theory*. *Minerva* 1: 54-74.
- Pöyhönen, S. 2016. Value of cognitive diversity in science. *Synthese* 194 (11): 4519-4540
- Ruphy, S. 2005. Why metaphysical abstinence should prevail in the debate on reductionism. *International Studies in the Philosophy of Science* 19(2): 105-121.
- Soler, L., Zwart, S., Lynch, M., and Israel-Jost, V. 2014. *Science after the Practice Turn in the Philosophy, History, and Social Studies of Science*. Routledge.
- Strevens, M. 2003. The rule of the priority rule in science. *The Journal of Philosophy* 100: 55-79.

- Strevens, M. 2013. Herding and the quest for credit. *Journal of Economic Methodology* 20(1): 19-34.
- Vaesen, K. and Katzav, J. 2017. How much each researcher receive if competitive government research funding were distributed equally among researchers? *Plos One* 2(9): e0183967
- Viola, M. 2015. Some remarks on the division of cognitive labor. *Roar Transactions*. 1: 1-14.
- Viola, M. 2018. Social epistemology at works: from Philosophical Theory to Policy Advice. *Roar Transactions* 6(1).
- Weisberg, M. and Muldoon, R. 2009. Epistemic landscapes and the division of cognitive labor. *Philosophy of Science* 76(2): 225-252.
- Wilholt, T. and Glimell, H. 2011. Conditions of Science: The Three-Way Tension of Freedom, Accountability and Utility. In *Science in the Context of Application*. Boston Studies in the Philosophy of Science.
- Wray, K.-B. 2000. Invisible hands and the success of science. *Philosophy of Science* 67(1): 163-175.
- Woody, A-I. 2014. Chemistry's Periodic Law: Rethinking Representation and Explanation After the Turn to Practice. In *Science after the Practice Turn in the Philosophy, History, and Social Studies of Science*, ed. Soler, L., Zwart, S., Lynch, M., and Israel-Jost, V. Routledge.
- Ylikoski, P. 1995. The invisible hand and science. *Science Studies* 8(2): 32-43.
- Ylikoski, P. and Aydinonat, N-E. 2014. Understanding with theoretical models. *Journal of Economic Methodology* 21(1): 19-36.
- Zollman, K.-J.-S. 2010. The epistemic benefit of transient diversity. *Erkenntnis* 72: 17-35.
- Zollman, K-J-S. 2018. The credit economy and the economic rationality of science. *The Journal of Philosophy* 115(1): 5-33.