

Natural Kinds and Concepts: A Pragmatist and Methodologically Naturalistic Account

Ingo Brigandt

Abstract: In this chapter I lay out a notion of philosophical naturalism that aligns with pragmatism. It is developed and illustrated by a presentation of my views on natural kinds and my theory of concepts. Both accounts reflect a methodological naturalism and are defended not by way of metaphysical considerations, but in terms of their philosophical fruitfulness. A core theme is that the epistemic interests of scientists have to be taken into account by any naturalistic philosophy of science in general, and any account of natural kinds and scientific concepts in particular. I conclude with general methodological remarks on how to develop and defend philosophical notions without using intuitions.

The central aim of this essay is to put forward a notion of naturalism that broadly aligns with pragmatism. I do so by outlining my views on natural kinds and my account of concepts, which I have defended in recent publications (Brigandt, 2009, 2010b). Philosophical accounts of both natural kinds and concepts are usually taken to be *metaphysical* endeavours, which attempt to develop a theory of the nature of natural kinds (as objectively existing entities of the world) or of the nature of concepts (as objectively existing mental entities). However, I shall argue that any account of natural kinds or concepts must answer to *epistemological* questions as well and will offer a simultaneously pragmatist and naturalistic defence of my views on natural kinds and concepts.

Many philosophers conceive of naturalism as a primarily metaphysical doctrine, such as a commitment to a physicalist ontology or the idea that humans and their intellectual and moral capacities are a part of nature. Sometimes such legitimate views motivate a more contentious philosophical program that maintains that any philosophical notion ought to be defined in a purely physicalist vocabulary (e.g., by putting forward a theory of concepts and intentional states that does not define them in terms of intentional notions). We will see that I reject this latter project (which is naturalistic in some sense) on naturalistic grounds, as science does not aim at developing reductive definitions. Rather than naturalism as a metaphysical doctrine, more germane to my account is a *methodological type of naturalism*. Here the idea is that some aspects of scientific method and practice should be used by philosophers in their attempts to develop philosophical accounts. I will illustrate this naturalistic method by laying out how philosophers can and ought to develop philosophical notions without simply relying on their personal intuitions or folk intuitions as revealed by experimental philosophy surveys. The starting point of this method is that philoso-

philical concepts are introduced for specific philosophical purposes – just like scientific concepts are used for particular epistemic purposes – and that these purposes determine the appropriateness of a philosophical analysis. For example, I will defend my account of what a ‘concept’ is in terms of its fruitfulness for explaining phenomena of interest to philosophers. My methodological naturalism aligns with a *pragmatist* perspective, since science is to a large extent a pragmatic enterprise and since scientists freely invoke scientific values and interests and defend their accounts in terms of their fruitfulness at meeting scientific aims.

Natural kinds

The traditional philosophical aim of an account of natural kinds is to provide a metaphysical characterisation of what a natural kind is, which among other things distinguishes natural kinds from other kinds and entities (Bird and Tobin, 2009). Given that, broadly speaking, a natural kind is a grouping of objects that corresponds to the objective structure of nature, an account of natural kinds must explain how natural kinds differ from nominal kinds, i.e., a grouping of various objects that is merely the result of human convention. One possibility is to define natural kinds as those kinds that figure in laws of nature (Fodor, 1974). Another, though not necessarily conflicting account may construe a natural kind as characterised by an essence, i.e., some intrinsic, structural property that all kind members share and that causes the typical properties associated with the kind (Putnam, 1975). For instance, the essence of oxygen is its atomic structure, including the number of subatomic particles making up an oxygen atom. This atomic structure explains in which chemical reactions oxygen can participate, and other chemical properties characteristic of the natural kind oxygen.

Functional kinds are generally considered to not be natural kinds; for if kind membership is defined by members having a common function, such a kind is bound to be structurally heterogeneous. This is because any one function is multiply realizable, i.e., there are different actual or possible physical entities that realize this function in one way or another. For instance, instances of ‘money’ as a functional kind include different metals (gold, coins), different objects made of paper (bills, cheques), and electronic states and information (computerized bank accounts). A functional kind from ecology such as ‘predator’ is likewise multiply realized. There is a plethora of predatory species across the animal kingdom that differ substantially in their structural-anatomical, physiological, developmental, and even behavioural properties. The structural heterogeneity of functional kinds is usually considered to be incompatible with them being natural kinds, which are deemed to be characterized by a shared structural essence.

Yet a look at biology suggests that many natural kinds are heterogeneous. A case in point is species and higher taxa, which have been taken to be prime examples of natural kinds in biology. The different individuals forming a biological species (e.g., orangutans) or a higher taxon (e.g., vertebrates) can be very different from each others. This variation is a biological reality and of scientific importance (e.g., underwriting the ability of species to evolve), so heterogeneity need not be an accidental feature, but can be constitutive of some kinds (Wilson et al., 2007). A good metaphysical solution is to make use of Richard Boyd's (1991, 1999a, 1999b) suggestion that natural kinds are *homeostatic property clusters* (HPC kinds). The identity of an HPC kind is, in general, not determined by a single essential property; instead, there is a cluster of properties that are correlated. Most of the kind members possess most of these properties, but none of the properties in the cluster has to be shared by all kind members, permitting variation among the members of an HPC natural kind. A requirement is that the correlation of properties is not an accident, but due to some mechanisms that causally maintain the correlation (thus the label 'homeostatic' property cluster). Thereby our grouping of objects into a kind based on such a cluster of correlated properties conforms to features in nature and HPC kinds are indeed natural rather than nominal kinds (Wilson et al., 2007).

Given the possibility of heterogeneous natural kinds, one may wonder whether at least some functional kinds are natural kinds after all. One option is to try to assess how many properties are correlated in a given kind, and count those kinds as natural where a sufficient number of properties are correlated. However, this approach ultimately results in a continuum between nominal kinds and natural kinds. More importantly, simply counting the number of properties correlated (or assessing the degree of homogeneity) fails to pay attention to the crucial *epistemic* role that natural kinds have for science. Natural kinds are important because they permit induction and explanation. In his discussion of the problem of induction, Nelson Goodman (1955) prominently argued that reliable induction requires projectible predicates. Within a contemporary, realist framework (which I adopt), it is clear that a predicate referring to a natural kind is projectible. The reliable correlation of properties in an HPC kind, for instance, grounds induction and other instances of scientific inference. Many natural kinds support scientific explanations, e.g., if the kind figures in laws or if the kind's essence (or one of the properties from the cluster defining an HPC kind) causes some of the features typically associated with a kind.

As I have argued in a previous paper, illustrated by biological examples (Brigandt, 2009), the above considerations suggest that the main philosophical task is not to offer a *metaphysical* construal of what a natural kind is, but an *epistemological* study of (i) what inferential and explanatory aims scientists pursue with the study of a certain natural kind, and (ii) how well a grouping of ob-

jects into a kind meets such inferential and explanatory aims.¹ For also many functional kinds studied for instance in ecology, physiology, psychology, or economics figure in scientifically important generalizations and explanations. While ‘money’ is multiply realized, it is part of many macroeconomic generalizations, such as Gresham’s law. Ecological generalizations exist for the temporal change of the sizes of predator and prey populations. Such generalizations do not describe internal, structural aspects of the members of a functional kind (e.g., organisms from different predator species), instead, the generalization pertains to *relations* between members of a functional and other kinds (e.g., predators from one species and their prey species). While the essences of natural kinds have typically been assumed to be intrinsic properties, relational properties are also important for many natural kinds from biology. A higher taxon (e.g., vertebrates) is defined as consisting of those organisms that are descended from a specific ancestral species, and ‘being descended from’ is a relational property. Likewise, many species concepts define a species in terms of relational properties, such as the ability to interbreed with other species members. This is why the HPC construal of kinds explicitly permits the cluster of properties that characterizes a natural kind to include both intrinsic and relational properties (Brigandt, 2009). Many biological entities (e.g., in molecular biology and physiology) possess their causal capacities only in certain contexts, so that biological explanations often involve relational properties (Brigandt, in press). Internal structure and structural homogeneity is not at all the only way to ground generalizations and explanations.

Many functional kinds support scientific generalizations and explanations, and thereby meet an epistemic hallmark of natural kinds. Debating whether or not some functional kinds really are natural kinds – based on some metaphysical construal of natural kinds – is moot. For any kind, the philosophically relevant question is an epistemic issue: how scientifically important is the grouping of an object into a kind, i.e., what generalizations and explanations can the kind figure in, and how important are they? This cannot be assessed in terms of how ‘natural’ or ‘real’ a grouping of objects is, or by simply counting how many properties are co-instantiated in a kind. Boyd (1999a, 1999b) – who I follow on this issue – has already highlighted the epistemic dimension of natural kinds by his notion of ‘accommodation’. The idea is that there are certain inductive and ex-

1 Even a traditional essentialist construal of natural kinds hints at this. An essence has two functions: (1) It determines kind membership (exactly those objects possessing the essence are kind members), and (2) it accounts for the kind’s characteristic properties (e.g., by virtue of the essence causing those properties). But condition (1) alone cannot define the notion of a natural kind, as even the members of a nominal kind share a property, namely, being a member of this conventionally defined kind. For a property metaphysically determining kind membership to be the essence of a natural kind, condition (2) has to apply, which points at the epistemic issue of whether this property is causally relevant.

planatory demands in science, and a grouping of objects is a natural kind to the extent to which these objects possess properties that fulfil – accommodate – the inductive and explanatory demands. Boyd's approach is naturalistic in that he puts forward a revised notion of 'natural kind' with the aim of capturing natural kinds as they are found in different scientific fields, including the special sciences. While the traditional construal of natural kinds as characterized by an intrinsic, structural essence applies to physico-chemical kinds only, kinds studied in biology, psychology, and the social sciences are nonetheless epistemically important by figuring in generalizations and explanations.

In summary, I do not think that there is any clear-cut metaphysical boundary between natural kinds and other kinds, and furthermore, I think that a purely metaphysical construal of what a natural kind is is of very limited use. The important philosophical task is an epistemological account of the role and relevance of various kinds in scientific theorizing. My account is *naturalistic* in that my only overarching philosophical theory is that natural kinds have to be philosophically studied based on the empirical details pertaining to each kind. These empirical considerations may differ from case to case. Some kinds have traditional essences, others are HPC kinds; some kinds figure in inductions only, others support explanations. Even within a scientific subdiscipline such as molecular biology, there are both kinds characterized by intrinsic properties, while other kinds are defined by relational properties. Most importantly, the empirical considerations pertaining to a kind include not only empirical properties of the kind's members, but also the *epistemic-scientific aims* that can be met by using the kind in scientific theorizing. These features matter for scientific practice and theory, so that a naturalistic approach must take them into account. The philosophical task ought to be a study of (i) what inferential and explanatory aims scientists pursue with the study of a certain natural kind, and (ii) how well a grouping of objects into a kind meets such inferential and explanatory aims (Brigandt, 2009; see also Love, 2009).

In addition to being naturalistic, this account of natural kinds is *pragmatist* as it essentially includes the goals of intellectual activity. I adopt a realism rather than an anti-realism about kinds in that it depends on the actual structure of the world whether the properties of kind members permit successful inference and explanation. At the same time, whether a kind is scientifically important (deemed a 'natural' kind) depends on what our inferential and explanatory interests are. The ineliminable role of human interests and values has been a core theme for several 19th and especially 20th century pragmatists; and I have argued that intellectual interests are relevant even for an account of natural kinds, which has traditionally been considered a purely metaphysical issue. I do not think that there is a unique representation of or classification scheme for the world, for there are various theoretical or intellectual goals that we (not nature) have, and different classifications or groupings into kinds may be needed to meet different

goals (Dupré, 1993). My account of natural kinds is not defended by reference to metaphysical considerations (such as a clear and principled distinction between natural kinds and other kinds); rather, I defend it in terms of its *fruitfulness* for understanding how natural kinds figure in scientific theorizing – a genuinely philosophical question.

Scientific concepts

This section lays out a theory of concepts, which I originally developed in recent writings (Brigandt, 2006, 2010b). The main motivation of this theory is to account for the rationality of semantic change, and I illustrate my framework by applying it to two cases of conceptual change in biology – the homology concept and the gene concept. A later section will discuss how I defend this account of concepts and draw implications for naturalism and pragmatism.

On my account, a scientific concept consists of three components of content: (1) the concept's reference, (2) the concept's inferential role, and (3) the epistemic goal pursued by the concept's use. There are two reasons for recognizing these three components. First, the different components of content (or different semantic properties of scientific terms) are ascribed for and fulfil different philosophical functions. Second, in the course of history a scientific concept may change in any of these components (and one component can change without the others). It should be obvious why I follow the tradition in considering reference to be part of a concept's content. Since two coreferential concepts may have a different epistemic role, philosophers often recognize a concept's sense or intension in addition to its reference. My version of this additional component is the concept's inferential role (as defended by inferential role semantics, also called conceptual role semantics). A term's inferential role is the set of inferences and explanations in which the term figures and which it supports in virtue of its specific content. The inferential role broadly aligns with the definition of a scientific term.

Central biological concepts may change in the course of history and their use may vary across different scientists, which is more precisely change or variation in a concept's inferential role and possibly its reference. I introduce the novel notion of the epistemic goal of a concept precisely because it accounts for the *rationality of semantic change and variation*. It is well-known that scientists pursue various epistemic goals; scientists aim at discovering different phenomena, making scientific inferences and confirming generalizations, and explaining various processes. A particular epistemic goal (e.g. explaining cell-cell communication) is often specific to a certain scientific field, in that it is pursued by this field (or a class of related fields), while other fields pursue other epistemic goals. Typically, many scientific concepts are deployed to pursue a given epis-

temic goal. My point here is that there are cases where an epistemic goal is tied to a specific scientific concept, in that the very rationale of introducing this concept and of continuing to use it is to pursue the epistemic goal. For example, the concept of natural selection is used to account for evolutionary adaptation. Some biological concepts are not used for any explanatory or deep theoretical purposes, but for the epistemic goal of discovering certain phenomena, as with some concepts from molecular and experimental biology. The below accounts of the homology concept and the gene concept will illustrate this idea in concrete cases. In a nutshell, the epistemic goal pursued by a scientific concept's use is the type of knowledge (certain kinds of inferences, explanations, discoveries) the concept is *intended* to deliver, given its usage by a research community.² (The inferential role, in contrast, is the set of inferences and explanations that the concept currently actually supports.)

A concept – or more precisely, its inferential role – embodies beliefs about the concept's referent. While scientists constantly acquire novel beliefs or discard previous ideas about the term's referent, these revised beliefs usually do not lead to a redefinition of the term under consideration. Thus, what has to be accounted for in the case of *semantic* change is why certain novel beliefs about a term's referent (but not others) warranted a change in the very inferential role (sense) of the term. On my account, the *epistemic goal* pursued by a term's use sets the standards for which possible changes in the term's inferential role count as rational. For instance, a concept's epistemic goal may be to explain certain phenomena, yet presently the concept's definition – reflecting available empirical beliefs – does not yield an adequate explanation of this kind. Once appropriate empirical insights become available, the concept's definition (inferential role) is revised, and this semantic change is warranted if in virtue of the new definition the concept supports the desired explanation. In general terms, change in a term's inferential role is rational if the new inferential role meets the term's epistemic goal to a higher degree than the term's prior inferential role. If rational change in inferential role entails a change in the term's reference, the latter is also rational. Sometimes, within a scientific discipline there is some variation even in the epistemic goal for which a term is used. This variation in the purposes for which a concept is used accounts for semantic variation, i.e., variation in the concept's inferential role and possibly its reference. While I introduce the notion of epistemic goal to account for the rationality of semantic change and variation, the traditional conceptual components of reference and inferential role are needed for a different purpose, namely, to account for how concepts make successful practice (verbal behaviour and interaction with the world) possible.

2 I do not maintain that an epistemic goal can be assigned to every scientific concept. As the notion of epistemic goal is to account for semantic change and variation (which usually only major theoretical concepts exhibit), it is sufficient to ascribe it to those concepts where semantic change or variation occur.

The change of the homology concept

Homology is a concept central to the practice of comparative and evolutionary biology (Brigandt, 2006; Brigandt and Griffiths, 2007). Homologous structures are the corresponding structures in different species. For example, the right arm in humans, the right forelimb in horses, the right flipper in dolphins, and the right wing of bats are homologous. Even the individual bones of the right forelimb (e.g., radius and ulna) reoccur in different species. Homologous structures are considered the same structures and given the same name in different species. Apart from bones, all sorts of anatomical structures can be homologous, such as organs, individual muscles, nerves, and tissues. Cell types and molecular structures such as genes are also considered homologous across species.

Homologous structures are present in different species due to inheritance from the species' common ancestor. This yields the post-Darwinian definition of homology: two structures in different species are homologous if they are derived from the same structure in the ancestor. Despite the evolutionary nature of homology, the homology concept was introduced at the beginning of the 18th century and was already an important concept in comparative biology well before the advent of Darwin's theory of evolution. In this pre-Darwinian period different non-evolutionary accounts were given for homologous structures. One idea was that different species are governed by the same laws of development, resulting in corresponding structures in different species. Another account appealed to abstract geometric body plans (or possibly to blueprints in the mind of God), so that structures in actual species were defined to be homologous in case they corresponded to the same element in the abstract body plan. This change in the definition of homology is an instance of semantic change, and it raises the following worry: Does the Darwinian revolution amount to the replacement of the pre-Darwinian concept of homology by a different concept, so that the term 'homology' switched from one to another concept? Are the pre-Darwinian and post-Darwinian concepts of homology incommensurable (meaning incommensurability in the sense of Kuhn, 1962 and Feyerabend, 1962)? The same issue has been raised by philosophers in a related context, namely, the concept of 'species' (Beatty, 1986). Are the pre- and post-Darwinian accounts of the nature species two distinct concepts? This opens the (philosophically unsatisfactory) possibility of arguing that what Darwin's *Origin of Species* actually showed was that there are no species – as defined by the pre-Darwinian concept of species.

To return to the homology concept, some semantic change did occur with the advent of evolutionary theory. Specifically, what I refer to as the concept's inferential role changed. However, there was also an important element of conceptual continuity, for the epistemic goal pursued by the use of the homology concept did not shift in the transition to Darwinism. Already before the advent of evolutionary theory, biologists used the homology concept for two epistemic

purposes: (1) the systematic morphological description of several species, and (2) the taxonomic classification of species. For both epistemic goals, individuating anatomical structures in terms of homology proved to be very fruitful. Another possible individuation scheme is analogy, where analogous structures are structures having the same function. The wings of birds and insects are analogous, but not homologous. Homologous structures need not be analogous, as the above example of the mammalian forelimb (human arm, dolphin flipper, bat wing) shows that the function of a homologous structure can be very different in different species. The very origin of the notion of homology is the recognition that corresponding structures can be present in relatively unrelated taxonomic groups (such as reptiles and mammals, or even fish and mammals), even though these structures have quite different shapes and different functions in different species. Homology individuates structures by breaking down an organism into its natural anatomical units (what these units are is not always obvious, as bones that are separate in some species can be fused in others). Homology also relates structures across species by identifying them as the same ones. This first yields unified morphological descriptions, far more unified than other, earlier individuation schemes permit. Many anatomical and developmental descriptions that apply to a structure in one species also hold for the corresponding, homologous structure in other species. The comparative practice using the homology concept made possible a unified morphological account of the vertebrate skeleton already before the advent of Darwinian evolutionary theory (Owen, 1849). Second, pre-Darwinian taxonomists aimed at grouping species into higher taxa not in an arbitrary or artificial fashion, but in a manner that revealed the species' so-called natural affinities. It became clear that while analogies were similarities independent of taxonomic relatedness, homologies across species reflected their natural affinities and were thus to be used as a guide to taxonomic relatedness.

The advent of Darwinism did not lead to a change in what comparative biologists attempted to achieve when using the homology concept – the epistemic goals were still systematic morphological description and the classification of species. Biologists came to adopt the novel, phylogenetic definition of homology precisely because they realized that it permitted them to pursue the traditional epistemic goals in a better fashion. Homologous structures came to be construed as structures derived from an ancestral structure. Taxonomic groups came to be seen as branches of a phylogenetic tree. Among other things, this explained why homologous structures (but not analogous structures) were to be compared in the classification of different species. A phylogenetic understanding of homology permitted a better resolution of controversial claims about particular homologies. A theoretically more sound morphology based on phylogenetic principles led to more adequate and unified anatomical descriptions encompassing different species. (For more details on the history of the homology concept see Brigandt, 2006.) In my terminology, the change in the

homology concept's inferential goal was rational because it permitted biologists to meet the concept's epistemic goals to a larger extent. There is no need to consider the conceptual change brought about by Darwinism as resulting in incommensurability.³

In this fashion, the notion of a concept's epistemic goal accounts for the rationality of semantic change. In addition to the homology concept's traditional use in comparative and evolutionary biology, in the second half of the 20th century this concept came to be used in two novel disciplines – molecular biology and evolutionary developmental biology. As I have argued earlier (Brigandt, 2003), each of these two new fields came to use the homology concept for somewhat different epistemic goals. This resulted in semantic variation across fields and in conceptual divergence, where homology is construed differently in systematics/evolutionary biology, in molecular biology, and in evolutionary developmental biology. Thereby, the notion of a concept's epistemic goal also accounts for why semantic variation is present.

The change of the gene concept

The gene concept is another case that, when submitted to my account, nicely illustrates the advantages of my framework of concepts. It is of particular interest since in the course of history the reference of the term 'gene' has changed, and even now the reference of this term may vary across uses. Since I have discussed the gene concept in a previous paper (Brigandt, 2010b), at this point I merely sketch the most interesting results, so as to move on to the implications of my account for pragmatism and naturalism.

The gene concept originated at the beginning of the 20th century, and by the 1930s, what is now called the classical gene concept had been clearly established. On my account, the epistemic goal pursued by the use of the classical concept is the prediction of patterns of inheritance, and the concept's inferential role (embodying knowledge about how genes as entities tied to chromosomal locations figure in sexual reproduction) met this epistemic goal to a large extent. With the advent of molecular genetics, out of the classical gene concept grew the molecular gene concept. All three components of conceptual content changed in this process. The epistemic goal of the molecular gene concept is not

3 Some may wonder whether the term 'homology' as used by pre- and post-Darwinian biologists involves the same concept or different concepts, given that its definition changed. Since on my account a term has three semantic properties (reference, inferential role, epistemic goal) and can change in each of them, I do not think that there is a unique account of concept individuation. No matter whether this instance of semantic change is viewed as an enduring homology concept changing internally or as one concept giving rise to a distinct concept, the rationality of the change in the term's semantic properties has to be explained, in this case through a change in inferential role (Brigandt, 2010b).

the prediction of patterns of inheritance (a process between generations), but the explanation of how genes produce RNAs and proteins as their molecular products (a process going on within single cells). The inferential role of the molecular concept includes the idea that genes are DNA sequences with a specific structure, as these structural features explain how genes figure in molecular mechanisms so as to code for their product – as demanded by the concept’s epistemic goal. Since classical genes are defined in terms of the phenotypic effects (as shown in inheritance patterns) and molecular genes are defined as structural units coding for proteins, the classical and the molecular concept may offer different accounts of how many genes are at genetic regions with a complex organization, entailing that the reference of the term ‘gene’ changed in the transition from classical to molecular genetics. The above discussion laid out how to account for change in a concept’s inferential role and reference by appeal to the concept’s stable epistemic goal. While in the transition from the classical to the molecular concept the very epistemic goal pursued by the use of the term ‘gene’ changed, this instance of semantic change can still be counted as rational, as discussed in Brigandt (2010b).

While the molecular gene concept was well-established by the 1970s, it has undergone semantic change in the last few decades, prompted by novel findings in molecular genetics and genomics. Originally, it was assumed that all genes have the same structure (a stretch of DNA delineated by a start and stop codon and preceded by a promoter sequence), where one such structural unit codes for a single product and every gene product results from one such DNA unit. However, it turned out that genes form a structurally heterogeneous kind and that the relation between DNA elements and their products is many-many. This led to revised construals of what molecular genes are, resulting in a historical change of both the inferential role and reference of the molecular gene concept. At the same time, the molecular gene concept’s epistemic goal has been stable – the concept is still used to explain how genes code for their products. The new use of the molecular gene concept came about by those findings about gene structure that bear on gene function (i.e., coding for gene products). Thereby it was an instance of rational semantic change, as current construals of what genes are provide an improved account of how DNA elements code for gene products – meeting the molecular gene concept’s epistemic goal to a higher degree.

This semantic change in the last few decades has also led to a significant degree of semantic variation. Nowadays, different molecular biologists may use the term ‘gene’ differently. One DNA segment can produce many different types of products, and several non-contiguous DNA segments can be involved in the coding of a single product. Due to these, among many other, complexities, it is unclear whether a given DNA element is an independent gene, the mereological sum of several genes, or just a mereological part of a gene. Different molecular biologists may use different criteria to individuate genes and even prefer differ-

ent construals of what a gene is. Likewise, the reference of the term ‘gene’ can vary from context to context. This results from the fact that rather than forming a single structural kind, genes are best viewed as set of many overlapping structurally defined kinds. Different biologists may be interested in different aspects of gene structure or function, and consequently refer to one or the other of these kinds. This instance of semantic variation can be philosophically understood as follows. Even though there is a generic epistemic goal pursued by all molecular biologists (namely, to account for gene function), this generic epistemic goal can be spelled out differently by different researchers. In a particular context, a more specific epistemic goal is usually in play. For example, some researchers focus on proteins as the gene product of interest, while others focus on RNAs as another kind of gene product. If different biologists address one and the same complex genetic region with different investigative or explanatory goals in mind, different accounts can result. Since different specific epistemic goals are legitimate, a context-sensitive use and reference of the term ‘gene’ is needed and in fact justified, so that the notion of a concept’s epistemic goal accounts for the rationality of semantic variation in addition to semantic change (for more details, see Brigandt, 2010b).

Pragmatism and methodological naturalism

Now the close relation between my theory of concepts and my account of natural kinds should become clear. My framework of concepts assumes that each scientific concept consists of three components of content: (1) the concept’s reference, (2) the concept’s inferential role (sense), and (3) the epistemic goal pursued in the concept’s use. Traditional theories of concepts assume that a concept consists in certain beliefs about a referent, such as a definition in terms of analytic statements (or an intension, or an inferential role). In addition to this, my framework takes into account what concept users (e.g., scientists) attempt to achieve by using those definitions and in revising them. This feature is the epistemic goal pursued by the concept’s use; and I introduce it as an additional component of conceptual content precisely because it accounts for the rationality of semantic change and variation. The epistemic goal is the type of knowledge – certain kinds of inferences, explanations, discoveries – the concept is intended to deliver. In the case of natural kinds, I argued that rather than putting forward a metaphysical characterization of what a natural kind is and applying it to concrete cases, the relevant philosophical project is the epistemological study of (i) what inferential and explanatory aims scientists pursue with the study of a certain natural kind, and (ii) how well a grouping of objects into a kind meets such inferential and explanatory aims. The first part can be rephrased as a study of the particular epistemic goal. In fact, I urge shifting the philosophical focus

from natural kinds (as a metaphysical category) towards various natural kind *concepts*, including a study of the epistemic goals for which these natural kind concepts are used. Similar to the way in which the epistemic goal sets the standards for when a new definition of a term (semantic change) is warranted, the epistemic goal underlying the use of a natural kind concept permits a philosophical evaluation of the extent to which the knowledge about a kind (a grouping of certain objects) meets the given epistemic goal.

A pragmatist approach to scientific concepts

My theory of concepts is pragmatist in three respects. First, while scientists may be aware of what epistemic goal they pursue with the use of a certain term (especially if other scientists use the term for somewhat different purposes), I do not require that the epistemic goal is a belief explicitly held by these scientists. Instead, the epistemic goal is constituted *implicitly* by how a scientific community uses a term. The epistemic goal is present and efficacious by influencing how a scientific community reacts to novel findings, e.g., by redefining the concept and thus changing subsequent use. The idea that conceptual content is implicit in practice is a central feature of Robert Brandom's (1994) type of pragmatism. Second, the epistemic goal of concept is a *non-truth-conditional aspect of meaning*; at least it differs in two ways from the traditional truth-conditional features of meaning. In line with the previous point, while semanticists can use propositions to describe a concept's epistemic goal, such an epistemic goal need not be an explicit propositional belief entertained by an individual scientist possessing this concept. Rather, the concept's epistemic goal is constituted implicitly by the term's usage in scientific practice, in fact by how an overall community uses the term. More importantly, unlike truth-conditional features of meaning, the epistemic goal does not consist in a belief about *states of the world* – it is not even a desire concerning how aspects of the world studied by science should be like. Instead, it is a goal concerning *scientific knowledge*, or a desire as to what a scientific community should achieve. Thereby my theory of concepts includes features that go beyond the function that concepts have in *representation*, and is in line with some 20th century pragmatists who have stressed the role of knowledge and language independently of representing the world (Rorty, 1979; Brandom, 1994). Third, my framework of concepts highlights the relevance of epistemic goals, interests, and values in science. Epistemic goals also interact with other components of conceptual content: a given epistemic goal motivates and justifies change in inferential role and reference, and likewise, a change in scientific beliefs and inferential roles may transform the epistemic goals that scientists currently deem to be most important. The relevance of interests and values and the entanglement of facts and values have been a major theme in many pragmatist approaches of the 19th and 20th century.

In addition to the pragmatist ingredients in my account of concepts, I offer a pragmatist defence of this approach. Concepts are commonly viewed as objectively existing mental entities (or as abstract entities grasped by mental states). However, I do not defend my account of concepts as the right *metaphysical* theory of what these mental entities really are. Instead, I put forward my framework because of its *fruitfulness* for philosophically studying scientific concepts and explaining their use and their historical change. For instance, I have claimed that the epistemic goal pursued by a concept's use is a semantic property of a term – just like reference and inferential role (intension) are – on the grounds that it is needed for a *semantic task*, namely, accounting for the rationality of semantic change and variation. Some may wonder whether the epistemic goal is actually a semantic or rather a pragmatic aspect of term use. But drawing the semantic-pragmatic distinction in a certain way is less important than the need to include the epistemic goal in any study of scientific concepts.

In a similar vein, one could contend that an epistemic goal is not a semantic property (and thus not a component of a concept), but an epistemic property (merely tied to concepts), maybe on the grounds that accounting for conceptual change is an epistemological rather than genuinely semantic task. Jerry Fodor (1994, 2001) has used this basic line repeatedly in defending his atomist theory of concepts. On Fodor's account, reference is the only component of conceptual content, so he excludes even intension or inferential role. Of course he has to acknowledge that co-referential concepts may differ in their inferential roles or have different intensions, which is due to different modes of representing the referent. But Fodor argues that the latter are epistemic, not semantic aspects of concepts, so while related to concepts, these epistemic properties are not part of concepts. In principle, one could make this distinction, while at the same time studying semantic properties (narrowly construed) and epistemic properties of concepts together. Yet the important point is that Fodor invokes the semantic-epistemic distinction precisely so that he can ignore all aspects of concepts apart from reference. He has likewise made plain that he does not want to be bothered with considering how to account for semantic change (Fodor, 2000).

In contrast, my first justification for recognizing three different components of content is that all of them are needed for important philosophical purposes. One task is to explain how concepts make successful practice possible, including communication across individuals and the interaction between agents and the world. The notion of reference is important for this but cannot do this philosophical task alone; inferential roles also have to be ascribed to terms to account for successful practice. Another philosophical task is to account for semantic change and variation, which requires the notion of an epistemic goal. How concepts support successful practice and why conceptual change occurs clearly are philosophical questions – deciding whether they are semantic or epistemological questions merely yields a verbal dispute. Second, not only are different compo-

nents of conceptual content to be recognized because they fulfil different philosophical functions, but they have to be studied *together* as they influence each other. For instance, the epistemic goal of a concept determines how the concept's inferential role and reference changes in the light of new empirical beliefs, so that the behaviour of one component cannot be understood without considering the others. These two reasons for viewing all three properties as components of a concept appeal to the philosophical usefulness of such a framework of concepts. One philosophical question is what theory of reference, inferential role, and a concept's epistemic goal can be provided that explains what metaphysically determines what exactly these properties are for a particular concept. From my perspective, the more relevant issue is that the philosopher is able to ascribe a particular referent, an inferential role, and an epistemic goal to a concrete concept in a scientific context, and to defend this ascription in terms of its fruitfulness in understanding the use and change of concepts. In other words, I view philosophical notions such as 'concept' and 'natural kind' as tools – tools that philosophers develop for a certain philosophical purpose, and that have to be defended in terms of how fruitful they are for this purpose.

A methodological naturalism in philosophy, modelled on scientific practice

So far I have laid out the pragmatist aspects of my account of natural kinds and my theory of concepts. They are in fact motivated by a naturalism, but my type of philosophical naturalism needs cashing out (Papineau, 2009). Many naturalists will be uncomfortable with my inclusion of epistemic goals in a theory of concepts. This is because many naturalistic philosophers are committed to what I call reductive naturalism. *Reductive naturalism* is the idea that philosophers should attempt to reduce philosophical notions to some scientific vocabulary, such as a physicalist vocabulary. For instance, normativity ought to be explained in terms of non-normative notions. Intentional and semantic notions such as 'reference', 'meaning', and 'mental representation' are to be reduced to non-intentional notions. Reductive naturalists may favour a causal theory of reference because it is deemed to be able to explain how the intentional phenomenon of referring arises from causal relations. Teleosemantic theories of mental representation attempt to explain how the normativity associated with the distinction between correct and incorrect representations is due to some causal-historical physical process analogous to natural selection (Neander, 2009). Fodor's (1990) asymmetric dependency theory of content appeals to laws between mental representations and their referents. Needless to say, my account of concepts does not conform to reductive naturalism. I used the semantic notions of reference and inferential role without offering a theory of how to reduce them to non-semantic notions. Moreover, to these traditional notions I added the notion of the epistemic goal pursued by a concept's use, invoking even interests and values.

In fact, I view reductive naturalism as misguided from a properly naturalistic perspective. Reductive naturalism is motivated by metaphysical naturalism, i.e., a commitment to a physicalist ontology. I endorse the idea that humans and their intellectual and moral capacities are a part of nature. But reductive naturalism first makes the additional epistemic demand that philosophical *notions* are to be reductively explained and second assumes that the reduction base is some physicalist vocabulary. Regarding the first issue, as I will discuss below, scientists do not aim at reducing their notions. Concerning the second issue, the assumed physicalist or ‘scientific’ vocabulary is often just a philosopher’s fiction. Even though they are disallowed by reductive naturalists, normative notions are routinely used by scientists when discussing scientific standards, proper method, relevant problems, and the proper aims of their disciplines (Brigandt, in press). Psychologists also use representational and intentional concepts without desiring to reduce them. Furthermore, philosophical theories deemed as naturalistic use notions not found in science. For example, Fodor’s (1990) theory of concepts postulates psychological ‘laws’ (obtaining between mental representations and their referents), and his asymmetric dependency account requires that some of these laws are metaphysically more basic than others, in that breaking the former breaks the others but not vice versa. Thereby he invokes counterfactuals, in fact counternomologicals, without asking whether psychologists assume laws and whether scientists make counterfactual statements of the kind Fodor needs. The kind of naturalism I endorse and view as most relevant to philosophy is a type of *methodological naturalism* that recommends that philosophical method be modelled on scientific method, and that philosophers develop and defend their notions in analogy to how scientists do this. I now point to some insights that philosophers can get from scientific practice.

Philosophical debates about reductionism in biology have shown that while experimental strategies and explanations that are reductionist in a certain sense exist, a thorough epistemic reductionism does not hold for biology (Brigandt and Love, 2008). Many biological explanations in terms of mechanisms involve molecular entities, but rather than explaining exclusively with reference to molecular features, mechanistic explanations shed light on the interaction among entities on several levels of organization (Craver, 2005; Darden, 2005). In a similar vein, numerous biological explanations result from appropriately coordinating knowledge from different biological subdisciplines. As a result, rather than developing philosophical models of epistemic reduction, a better way of understanding scientific theorizing is to analyze the *integration* of knowledge and explanations across different biological fields and levels of organization (Brigandt, 2010a; Darden and Maull, 1977; Grantham, 2004; Love, 2008).

Alan Love and I have advocated a problem-based philosophical account of epistemic integration and interdisciplinary explanation (Love, 2008; Brigandt, 2010a; Brigandt and Love, 2010). Some scientific problems may motivate inte-

gration, in that they can only be solved by integrating items of knowledge from several fields. For instance, one of the core problems of current evolutionary biology is to account for the evolutionary origin of novel structures, such as the evolution of vertebrate jaws (the transition from jawless to jawed vertebrates), the origin of fins in fish, or the evolution of feathers in birds. Biologists are well aware of the fact that the explanation of novelty requires the integration of knowledge from evolutionary genetics, developmental biology, morphology, paleontology, phylogeny, and ecology, among others. Moreover, a problem such as explaining evolutionary novelty is complex in that it has an implicit structure consisting of different component questions and their hierarchically organized relations. Tied to this are criteria of explanatory adequacy that indicate acceptable answers to a component question (Love, 2008). The internal structure of a complex problem and the associated criteria of explanatory adequacy foreshadow how the contributions from different intellectual disciplines are to be coordinated. Therefore a scientific problem may not only motivate integration, but suggest which disciplines are to be involved and how their knowledge is likely to be related and integrated.

Recently I have argued that in explanations of the origin of evolutionary novelty, developmental biology is particularly crucial, where not only developmental genes, but also higher-level developmental entities and processes are important (Brigandt, 2010a). The reason is that organismal features on different levels of organization can evolve independently of each other, so that an important task is to understand the causal relations and partial developmental dissociations of entities on several levels that make independent evolution on different levels possible. As a result, there are scientific problems where disciplines studying higher level entities carry greater explanatory force. A common assumption is that the ontological hierarchy from smaller to larger entities yields an *epistemological* hierarchy of scientific fields (microphysics, chemistry, molecular biology, organismal biology, sociology), where a lower-level theory can reduce the ones on higher levels and thus is always explanatorily more fundamental. However, my considerations show that which fields are more fundamental in complex explanations varies with the problem pursued (Brigandt 2010a). Philip Kitcher (1999) maintains that while complete unification in science cannot actually be achieved given the complexity of nature, unification is a regulative ideal. In contrast, I have argued that unification/integration is neither a regulative ideal nor an aim in itself. Instead, a certain kind of integration may be needed *for the aim of* solving a scientific problem (Brigandt, 2010a). At the same time, solving a complex problem also requires a certain degree of scientific specialization, so that integration and interdisciplinarity often goes together with disciplinary specialization (Bechtel, 1986). Explanatory integration in biology is not the merging or the stable unification of different fields, but the transient coordination of knowledge from different fields solely for the purposes of a

specific problem, where different interdisciplinary intellectual relations obtain for different problems. This problem-based account of integration illustrates yet again the relevance of epistemic interests, goals, and values in science as a recurring pragmatist theme of this essay, and the philosophical need to take them into account in any attempt to understand the structure and dynamics of science.⁴ In a similar vein, I do not view *truth* as a regulative ideal in science. True representations are needed to meet various aims that scientists are interested in, but truth is not an aim in itself. Moreover, many scientific models involve various idealizations and thus are not literally true representations of the world. Not only are idealizations unavoidable, but some scientific models are explanatory (or have other epistemic virtues) precisely because they make those idealizations (Cartwright, 1983). And since different models of the same phenomenon make different idealizations that are jointly incompatible, it is often not possible to combine these models into an account that more completely represents nature (Kellert et al., 2006).

Let me return to the difference between reductive naturalism and my methodological naturalism. When philosophers talk about ‘explanation’ and attempt at giving explanations, they often have in mind a definition of a core philosophical term (e.g., ‘knows that *p*’) by means of necessary and sufficient conditions. For a reductive naturalist, such a definition has to be in terms of a physicalist (or alleged scientific) vocabulary. In contrast, ‘explanation’ in biology is something different. Biologists do not at all aim at a definition of ‘life’. Instead, the aim is to gain an understanding of the causal workings of various life phenomena. As we saw above, in such causal explanations there is no privileged level of explanation (e.g., the molecular level) and no privileged vocabulary (e.g., molecular terminology) in which all biological explanations should be given. Biological explanations usually account for the interrelation and interaction of entities on various levels. A methodologically naturalistic philosophy ought to proceed in the same way. In the context of philosophy of science, this means that naturalistic philosophers are not confined to merely describe what some call ‘science’, but may use normative concepts and distinguish between genuine science and pseudo-science just like scientists use normative notions to discuss proper method, well-confirmed hypotheses, and the proper aims of their disciplines. Rather than aiming at a definition of ‘knowledge’, naturalistic philosophers have to study the normative, cognitive, and social factors involved in *knowledge pro-*

4 While my discussion has focussed on epistemic interests in science, I do not think that a defensible distinction between epistemic and other values can be drawn. In current biomedical research interests and values that are intuitively epistemic and intuitively non-epistemic are so entangled that they have to be studied together and that grouping them into two classes is of no philosophical relevance. Major approaches in contemporary philosophy of science that are aware of the relevance of values in science are studies of the social dimensions of scientific knowledge and feminist philosophy of science.

duction – just as biologists do not define ‘life’ but study life processes. Rather than avoiding normative notions or reducing them to a non-normative vocabulary, the philosophical aim ought to be to shed light on the *interplay* of normative, cognitive, and social aspects in science – in analogy to biologists offering explanations describing the interaction of phenomena on different levels.

In the case of philosophy of mind and language, for a methodological naturalism modelled on scientific practice there is no *a priori* need to reduce intentional notions (e.g., ‘reference’ or ‘inferential role’) or normative notions (e.g., ‘correct mental representation’ or ‘a concept’s epistemic goal’). My account of concepts ascribes different intentional or normative properties to scientific terms, and defends such an ascription in concrete cases by the explanatory insights it yields for philosophy, namely accounting for how concepts support successful scientific practice and can undergo rational change. The focus, on my approach, is on studying the interrelations among the three components of content, for instance by laying out how in the course of history one component changes while others are stable, and how change in one component is made rational by another component of content. In sum, rather than reducing a philosophical concept to a physicalist vocabulary, the general philosophical aim should be to shed light on the relations and interactions among phenomena referred to by various philosophical concepts and phenomena studied by the natural, cognitive, and social sciences.

Philosophical method beyond intuitions and experimental philosophy

Philosophers often analyze philosophical concepts using their intuitions. The idea is that a philosophical account (e.g., of ‘knowledge’) should conform to intuitions as to how this concept applies to various cases and hypothetical scenarios. For instance, the intuitions elicited by Gettier cases were used to reject the traditional analysis of knowledge as justified true belief, and motivated alternative accounts of knowledge (Gettier, 1963). Kripke (1972) used his famous Gödel/Schmidt case to criticize former descriptive theories of reference and to argue for his causal theory of reference. Cases of causal pre-emption have been used to argue against analyses of causation in terms of counterfactuals (Lewis, 1973, 2000). In a similar vein, trolley cases have prominently been used to test normative ethical theories. A potential problem with this philosophical method is that there seem to be counterexamples to every analysis, and that no single account appears to be able to cover all cases. However, an analysis need not be in terms of a short and simple statement (e.g., ‘*S* knows that *p* iff *p* is true, *S* believes that *p*, and *S* is justified in believing that *p*’). It may well be possible that there is a very complex, conditional statement that covers all cases. Much more problematic is the fact that different people can have different intuitions about one and the same case (Mallon et al., 2009). If a disagreement about intuitions is

stable and the use of the reflective equilibrium method leads to two rival accounts where one fits the intuitions of philosopher *A* and the other fits the intuitions of philosopher *B* to the same degree, it is unclear what would settle such a conflict of intuitions.

In my view, the real problem with philosophers' reliance on intuitions is this: all that can be achieved by the method of considering various cases that illicit one's intuition about how a philosophical concept applies to these cases is to reveal one's *current* conception. The philosophical aim, however, should not be to explicate one's current concept of, say, knowledge, but to put forward an adequate account of knowledge. An intuition-based method tacitly assumes that we have already understood the philosophical phenomenon referred to by the concept and just have to make explicit the concept grasped. This is false. Our current concepts are likely to be impoverished and in need of improvement, as shown by the history of scientific concepts. The pre-Darwinian concept of homology (offering a non-evolutionary definition of homology) enabled a relatively successful practice, but embodied theoretical misconceptions that were cleared up with the advent of the post-Darwinian concept of homology so as to permit further practical progress in comparative biology. The molecular gene concept (and its novel way to define what genes are) yielded a significant advance over the classical gene concept. Philosophers should likewise strive to improve their concepts rather than merely explicating their current conceptions.

The recent area of *experimental philosophy* is opposed to armchair philosophy and deemed a naturalistic approach, as its method of designing questionnaires and statistically evaluating responses mimics the method of psychology. Experimental philosophers have chartered the between-person variation in intuitions about one and the same hypothetical case, including laypersons in their studies (Machery et al., 2004). However, if a single person's armchair intuitions are philosophically irrelevant – as I have alleged – then surveying the intuitions of many persons is equally useless, to the extent that it just describes various current conceptions rather than developing an improved account. Experimental philosophy is interested in understanding the cognitive processes that generate intuitions about philosophical issues. Some may view this as a means to an end: we are ultimately interested in the metaphysical phenomenon of causation, and use knowledge about cognition (how intuitions about causation are generated) for this metaphysical purpose. In contrast, the experimental philosophers Knobe and Nichols (2008) maintain that understanding the psychological mechanisms underlying philosophical intuitions is an end in itself. If so, then experimental philosophy pursues an aim different from traditional philosophy and is not interested in developing accounts of causation, reference, knowledge, etc. Other experimental philosophers assume that knowledge about cognitive processing sheds light on the reliability of philosophical intuitions, for instance, if emotional and affective factors influence folk intuitions about intentional agency

(Nadelhoffer, 2006). But in line with my above point, even if we knew which intuitions held by philosophers were reliable, we would still have nothing but a description of our current conception. The results of experimental philosophy about the unreliability or between-person variation of philosophical intuitions can indeed be used to reject previous philosophical accounts based on intuitions. But this would just be a critique of earlier accounts and fall short of the constructive task of developing an improved account.

The approach motivated by my science-based methodological naturalism goes beyond both armchair philosophy and experimental philosophy. Scientists modify the definitions of their concepts, and the aim is not to explicate current conceptions, but to arrive at a revised account. My theory of scientific concepts explains what makes semantic change rational, namely, the epistemic goal pursued by a concept's use sets the standard for when a change in the concept's definition is progressive. My methodologically naturalistic strategy is to recommend the same approach for *philosophical* concepts. Philosophical concepts such as 'knowledge' are not just lay notions that we have already grasped and implicitly understood, but they are technical terms introduced for certain philosophical purposes. A concept is an intellectual tool used in the attempt to meet a specific intellectual goal. Any philosophical account is not to be judged in terms of how well it conforms to intuitions, but how well it actually meets the respective *philosophical goal*.⁵

For example, the goal of an account of knowledge is to understand particular aspects of doxastic states. Contrastive accounts of knowledge analyze it as a three-place relation: '*S* knows *p* rather than *Q*', where *Q* is a set of alternative propositions (Karjalainen and Morton, 2003). A common objection is that this is not 'what we mean' by knowledge. This may be true but is philosophically irrelevant. The relevant question is whether a contrastive account of knowledge

5 David Papineau (this volume) also favours a methodological naturalism. However, he views his account not as a revision but as a description of current philosophical practice, claiming that even armchair philosophy is already naturalistic (and that those who say otherwise misdescribe their practice). Papineau argues that all interesting philosophical claims are synthetic statements, and suggests that armchair reflection can support synthetic philosophical claims just as thought experiments in physics yield synthetic claims. The problem with this analogy is that Papineau fails to realize that thought experiments in physics are just arguments (Norton 1996), so that they yield synthetic claims from synthetic premises – which must (and can) be justified by experience. Defenders of armchair philosophy may deny that the premises of philosophical arguments are essentially contingent upon experience. In any case, Papineau's thought experiment analogy fails to explain how philosophers can justify their claims. My methodological naturalism is focused on improving current philosophical practice, by showing how philosophers can arrive at and justify their accounts. Namely, the philosophical goal of an account (e.g., an account of 'reference') entails what would count as an improved account and what empirical information (*a posteriori* knowledge) is germane to the defence of such an account.

has resources for understanding doxastic states that a traditional two-place analysis of knowledge does not have. To take to concept of ‘reference’, here a philosophical goal is to account for an agent’s verbal and intentional behaviour including interactions with objects. To be sure, ordinary persons are already in the business of interpreting each other, so that their assumptions about who refers to what entities with a certain word are likely to be true. However, such assumptions are philosophically relevant not because they are common intuitions about reference, but because they are effective for the task of semantic interpretation and thus contribute to meeting the philosophical goal of an account of reference. Experimental philosophy can assemble detailed data on what intuitions different persons have about different issues and how they use philosophical terms. But without a clearly spelled out philosophical goal, it is unclear which of this data is relevant, and how to use it to develop and defend a philosophical account. As the example of ‘reference’ shows, *empirical* data is part of such a justification, so that philosophy is not an *a priori* business. For whether a certain general account of reference (or an assignment of referents to particular terms) is successful depends on how well it interprets the verbal behaviour of agents, which (apart from normative considerations) involves empirical facts of how persons actually behave toward each other and toward objects.

It should now be clear that we have already seen my methodological naturalism about developing philosophical concepts in action. For my concept of a ‘concept’ was defended precisely in these terms. I did not argue for my theory of concepts by claiming that it conformed to intuitions about concepts (such as the idea that concepts are about intentionality, so that only reference matters to concepts). Rather, I laid out two philosophical goals that any account should meet: (1) to understand how concepts support successful practice, including communication and interaction with the world, and (2) to understand how concepts can rationally change and vary across persons. My contention was that my framework is fruitful for this purpose, and that all three components of conceptual content (reference, inferential role, and epistemic goal) are needed for these two aims. This may be contended, but only by putting forward an account of concepts that is better at meeting the two philosophical goals (or by arguing that there is a further relevant goal that my account neglects). Thereby I have approached the philosophical concept of a ‘concept’ just like scientists develop, defend, and improve scientific concepts.

From this science-based methodologically naturalistic perspective, the first step in any philosophical analysis should be to get clear about the philosophical goal pursued, and subsequently to clarify the standards that determine the adequacy of any analysis relative to this goal. In attempts to meet such a philosophical goal, three different issues can turn out to be the case. First, it may be revealed that a contextualist or pluralist account is needed to meet the respective goal. For instance, if no simple and monolithic definition of ‘knowledge’ can be

satisfactory, a contextualist account of knowledge may be the solution (Rysiew, 2008). In the case of theories of 'reference', for some complex scientific concepts such as the gene concept it may be impossible to assign a unique referent to some utterances of the term 'gene'. Different assignments of referents may work equally well to account for successful communication and verbal interaction with the world, so that a pluralist theory of reference can be called for (Burian et al., 1996) and does not result in semantic relativism. I view the concept of 'causation' as a technical term that philosophers use to understand causal explanation in science. Even if physics does not use the concept 'causation' at all (as it has better notions to account for the dynamics of physical processes), a philosophical account of causation can be fruitful by shedding light on causal explanation in other scientific disciplines apart from physics. Given the different kinds of causal explanations, a contextualist or pluralist account of causation is likely to be needed. Second, it may turn out that one philosophical concept alone cannot effectively meet the philosophical goal, and instead that one concept has to be split into several concepts, which are to be distinguished, or that novel concepts have to be introduced. Possibly both a concept of 'knowledge' in terms of reliably tracking truth and a concept of 'knowledge' as beliefs reflectively defensible by the believer may capture some epistemically important feature of doxastic states. Paying attention to the philosophical goal in play permits a justification for using a contextualist account or for introducing and distinguishing different concepts. Finally, it may turn out that what was taken to be one philosophical goal is better divided into several related goals.

Conclusion

My type of naturalism is a methodological naturalism modelled on scientific practice. The general tenet is that philosophers can benefit from emulating some aspects of scientific practice; in fact, they should develop and defend philosophical notions in analogy to how scientists do this. For instance, biology offers explanations that do not reduce phenomena to the molecular level but account for the interaction among phenomena on several levels of organization. Since such interdisciplinary explanations are very successful, philosophy can follow this model, so that philosophical concepts do not have to be reduced to a non-normative or even physicalist vocabulary. Instead, philosophers should study the relations and interactions among phenomena referred to by various philosophical concepts and phenomena investigated by the natural, cognitive, and social sciences. Furthermore, scientists do not put forward scientific concepts by using their intuitions. Rather than explicating their current understanding of a concept, scientists strive at putting forward an improved account, where this revised account is defended in terms of its scientific fruitfulness. In analogy to scientific

concepts being used to pursue specific epistemic goals, philosophical concepts ought to be viewed as technical terms introduced for certain philosophical purposes. A philosophical concept is an intellectual tool and to be improved and defended in terms of its fruitfulness for meeting a philosophical goal.

Since science is to a large extent a pragmatic enterprise, such a methodologically naturalistic approach aligns with a pragmatist perspective for two reasons. First, a philosophical account of empirical entities (e.g., natural kinds, concepts) involves scientific notions (e.g., laws, mechanisms) and how science studies those entities (e.g., the aims of scientific explanations), so that philosophical accounts have to do justice to and reflect actual science and its practice. Second, if putting forward philosophical notions (e.g., natural kind, concept) is analogous to scientific notions being put forward, then philosophers' practice and their ways of supporting their accounts ought to be as pragmatist as science is, including paying attention to one's intellectual aims and interests and defending one's concepts as fruitful tools.⁶

University of Alberta

References

- Beatty, J. (1986) 'Speaking of species: Darwin's strategy' in D. Kohn, ed., *Darwin's Heritage: A Centennial Retrospect*, Princeton: Princeton University Press.
- Bechtel, W. (1986) 'Integrating sciences by creating new disciplines: The case of cell biology' *Biology and Philosophy*, 8, 277-299.
- Bird, A. & Tobin, E. (2009) 'Natural kinds' in E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Spring 2009 Edition.
- Boyd, R. (1991) 'Realism, anti-foundationalism and the enthusiasm for natural kinds' *Philosophical Studies*, 61, 127-148.
- (1999a) 'Homeostasis, species, and higher taxa' in R. A. Wilson, ed., *Species: New Interdisciplinary Essays*, Cambridge, MA: MIT Press.
- (1999b) 'Kinds as the 'workmanship of men': Realism, constructivism, and natural kinds' in J. Nida-Rümelin, ed., *Rationality, Realism, Revision: Proceedings of the 3rd International Congress of the Society for Analytic Philosophy*, Berlin: de Gruyter.
- Brandom, R. B. (1994) *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Cambridge, MA: Harvard University Press.
- Brigandt, I. (2003) 'Homology in comparative, molecular, and evolutionary developmental biology: The radiation of a concept' *Journal of Experimental Zoology (Molecular and Developmental Evolution)*, 299B, 9-17.

6 I am indebted to Catherine Clune-Taylor and Jonathan Knowles for comments on a draft of this essay. I thank Jonathan Knowles and Henrik Rydenfelt for organizing the publication of this volume. The work on this essay was supported by the Social Sciences and Humanities Research Council of Canada (Standard Research Grant 410-2008-0400).

- (2006) *A Theory of Conceptual Advance: Explaining Conceptual Change in Evolutionary, Molecular, and Evolutionary Developmental Biology*, Dissertation, University of Pittsburgh.
- (2009) 'Natural kinds in evolution and systematics: Metaphysical and epistemological considerations' *Acta Biotheoretica*, 57, 77-97.
- (2010a) 'Beyond reduction and pluralism: Toward an epistemology of explanatory integration in biology' *Erkenntnis*, 73, 295-311.
- (2010b) 'The epistemic goal of a concept: Accounting for the rationality of semantic change and variation' *Synthese*, 177, 19-14.
- (in press) 'Philosophy of biology' in S. French & J. Saatsi, eds., *The Continuum Companion to the Philosophy of Science*, London: Continuum Press.
- Brigandt, I. & Griffiths, P. E. (2007) 'The importance of homology for biology and philosophy' *Biology and Philosophy*, 22, 633-641.
- Brigandt, I. & Love, A. C. (2008) 'Reductionism in biology' in E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Fall 2008 Edition.
- (2010) 'Evolutionary novelty and the evo-devo synthesis: Field notes' *Evolutionary Biology*, 37, 93-99.
- Burian, R. M., Richardson, R. C. & van der Steen, W. J. (1996) 'Against generality: Meaning in genetics and philosophy' *Studies in History and Philosophy of Science*, 27, 1-29.
- Cartwright, N. (1983) 'The truth doesn't explain much' in N. Cartwright, ed., *How the Laws of Physics Lie*, Oxford: Oxford University Press.
- Craver, C. F. (2005) 'Beyond reduction: Mechanisms, multifield integration and the unity of neuroscience' *Studies in the History and Philosophy of Biological and Biomedical Sciences*, 36, 373-395.
- Darden, L. (2005) 'Relations among fields: Mendelian, cytological and molecular mechanisms' *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 349-371.
- Darden, L. & Maull, N. (1977) 'Interfield theories' *Philosophy of Science*, 44, 43-64.
- Dupré, J. (1993) *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Cambridge, MA: Harvard University Press.
- Feyerabend, P. (1962) 'Explanation, reduction, and empiricism' in H. Feigl & G. Maxwell, eds., *Scientific Explanation, Space, and Time*, Minneapolis: University of Minnesota Press.
- Fodor, J. A. (1974) 'Special sciences (or: The disunity of sciences as a working hypothesis)' *Synthese*, 28, 97-115.
- (1990) 'A theory of content II: The theory' in J. A. Fodor, ed., *A Theory of Content and Other Essays*, Cambridge, MA: MIT Press.
- (1994) *The Elm and the Expert: Mentalese and its Semantics*, Cambridge, MA: MIT Press.
- (2000) 'Replies to critics' *Mind and Language*, 15, 350-374.
- (2001) 'Language, thought and compositionality' *Mind and Language*, 16, 1-15.
- Gettier, E. (1963) 'Is justified true belief knowledge?' *Analysis*, 23, 121-123.
- Goodman, N. (1955) *Fact, Fiction, and Forecast*, Cambridge, MA: Harvard University Press.
- Grantham, T. A. (2004) 'Conceptualizing the (dis)unity of science' *Philosophy of Science*, 71, 133-155.
- Karjalainen, A. & Morton, A. (2003) 'Contrastive knowledge' *Philosophical Explorations*, 6, 74 - 89.

- Kellert, S. H., Longino, H. E. & Water, C. K. (2006) 'Introduction: The pluralist stance' in S. H. Kellert, H. E. Longino & C. K. Water, eds., *Scientific Pluralism*, Minneapolis: University of Minnesota Press.
- Kitcher, P. (1999) 'Unification as a regulative ideal' *Perspectives on Science*, 7, 337-348.
- Knobe, J. & Nichols, S. (2008) 'An experimental philosophy manifesto' in J. Knobe & S. Nichols, eds., *Experimental Philosophy*, Oxford: Oxford University Press.
- Kripke, S. (1972) 'Naming and necessity' in D. Davidson & G. Harman, eds., *Semantics of Natural Language*, Dordrecht: Reidel.
- Kuhn, T. S. (1962) *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press.
- Lewis, D. (1973) 'Causation' *The Journal of Philosophy*, 70, 556-567.
- (2000) 'Causation as influence' *The Journal of Philosophy*, 97, 182-197.
- Love, A. C. (2008) 'Explaining evolutionary innovations and novelties: Criteria of explanatory adequacy and epistemological prerequisites' *Philosophy of Science*, 75, 874-886.
- (2009) 'Typology reconfigured: From the metaphysics of essentialism to the epistemology of representation' *Acta Biotheoretica*, 57, 51-75.
- Machery, E., Mallon, R., Nichols, S. & Stich, S. P. (2004) 'Semantics, cross-cultural style' *Cognition*, 92, B1-B12.
- Mallon, R., Machery, E., Nichols, S. & Stich, S. (2009) 'Against arguments from reference' *Philosophy and Phenomenological Research*, 79, 332-356.
- Nadelhoffer, T. (2006) 'Bad acts, blameworthy agents, and intentional actions: Some problems for juror impartiality' *Philosophical Explorations*, 9, 203-219.
- Neander, K. (2009) 'Teleological theories of mental content' in E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Winter 2009 Edition.
- Norton, J. D. (1996) 'Are thought experiments just what you thought?' *Canadian Journal of Philosophy*, 26, 333-366.
- Owen, R. (1849) *On the Nature of Limbs: A Discourse*, edited by R. Amundson (with a preface by B. K. Hall and introductory essays by R. Amundson, K. Padian, M. P. Winsor and J. Coggon), Chicago: University of Chicago Press, 2007.
- Papineau, D. (2009) 'Naturalism' in E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Spring 2009 Edition.
- (this volume) 'The poverty of analysis' (originally published in *Aristotelian Society Supplementary Volume*, 2009, 83, 1-30).
- Putnam, H. (1975) 'The meaning of 'meaning'' in H. Putnam, ed., *Mind, Language and Reality: Philosophical Papers, Vol. 2*, Cambridge: Cambridge University Press.
- Rorty, R. (1979) *Philosophy and the Mirror of Nature*, Princeton: Princeton University Press.
- Rysiew, P. (2008) 'Epistemic contextualism' in E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Spring 2009 Edition.
- Wilson, R. A., Barker, M. J. & Brigandt, I. (2007) 'When traditional essentialism fails: Biological natural kinds' *Philosophical Topics*, 35, 189-215.