

Where Do You Get Your Protein? Or: Biochemical Realization

Tuomas E. Tahko 

ABSTRACT

Biochemical kinds such as proteins pose interesting problems for philosophers of science, as they can be studied from the points of view of both biology and chemistry. The relationship between the biological functions of biochemical kinds and the microstructures that they are related to is the key question. This leads us to a more general discussion about ontological reductionism, microstructuralism, and multiple realization at the biology–chemistry interface. On the face of it, biochemical kinds seem to pose a challenge for ontological reductionism and hence motivate a dual theory of chemical and biological kinds, a type of pluralism about natural kinds. But it will be argued that the challenge, which is based on multiple realization, can be addressed. The upshot is that there are reasonable prospects for ontological reductionism about biochemical kinds, which corroborates natural kind monism.

- 1 *Introduction*
 - 2 *Functions: Aetiological or Goal-Directed?*
 - 3 *Moonlighting and Multiple Determinations*
 - 4 *The Powers-Based Subset Strategy*
 - 5 *The Case of Haemoglobin*
 - 6 *Haemoglobin and the Problem of Lower-Level Vengeance*
 - 7 *Multiple Realization and (Higher-Order) Interest Relativeness*
 - 8 *The Prospects for Ontological Reductionism*
-

1 Introduction

Philosophers of science have recently taken an interest in biochemical kinds, namely, natural kinds at the intersection of chemistry and biology, such as proteins (for example, Slater [2009]; Tobin [2010]; Goodwin [2011]; Bartol [2016]; Havstad [2018]). Proteins are macromolecules, which can be studied

from both biological and chemical perspectives.¹ If macromolecules are considered to be chemical kinds, we would typically explain their chemical properties in terms of their physical structure. This view of natural kinds is sometimes called ‘microstructuralism’ or ‘microstructural essentialism’, suggesting that microstructure is definitive of natural kinds and may to be used to classify them.² This is the chemical perspective concerning biochemical kinds. But when considered to be biological kinds, the role of these molecules in physiological processes is important. This role is typically called the (biological) function of the kind. We can talk about ‘functional kinds’ more generally, emphasizing that it is the function that we use as the basis for classification. The biological perspective also involves evolutionary or aetiological considerations, given that biological functions are the result of a causal sequence of evolutionary processes. This is the biological perspective concerning biochemical kinds.

The distinction between chemical and biological perspectives reflects some classic debates in philosophy of science. But it is not, of course, scientifically accurate: the various ‘levels’ of science are intertwined to such an extent that this type of clear division between biology and chemistry is certainly not available. Yet, as we will see, many anti-reductionist arguments (consider, for example, Fodor [1974]) start with the assumption that the special sciences do operate on distinct levels. So, I will adopt such a distinction as my working hypothesis, even though I will eventually propose that the distinction is based on a pragmatic rather than an ontological difference.

The key question to be examined in this article is the following:

Bio-Chem: What is the relationship between the biological function of a macromolecule and the microstructure related to that function?

A possible answer to this question can be seen in terms of identity-based reduction, in which case the biological function would be identical to the microstructure.³ This can be called ‘one-one reduction’ or ‘ontological

¹ In what follows I will continue to talk about the biological and the chemical perspective without any assumptions about which perspective practising biologists or chemists have. Indeed, it is likely that they consider both perspectives. But when it comes to the philosophical analysis of this interface, the perspective matters, as we will see.

² While microstructuralism about chemical kinds is not universally accepted, it is a popular view, largely because of the influence of Kripke and Putnam. The chemical perspective will be understood in terms of microstructuralism in what follows. However, microstructuralism does not necessarily entail that all the microstructural properties are also intrinsic. For discussion, see (Needham [2011]; Williams [2011]; Hendry [2012]; Tahko [2015]). The present article is concerned with the biology–chemistry interface and will be neutral about the chemistry–physics interface, even though, as an anonymous referee has pointed out, one may question the metaphysical significance of the debate about biochemical kinds depending on the status of chemical kinds themselves.

³ The term ‘microstructure’ is deliberately vague. It can refer either to the chemical structure of a chemical substance, in virtue of which the substance can undergo chemical reactions, or to the physical microstructure. I am interested in whatever the relevant lower-level structure related to

reduction'.⁴ There are two important challenges for ontological reductionism. The first is the possibility of one microstructure realizing multiple biological functions. This can be called 'one-many determination' or 'multiple determinations'.⁵ The second is the possibility of multiple microstructures realizing the same biological function. This is the more familiar case of 'multiple realization' or 'many-one realization'.⁶ I will use these labels interchangeably in what follows. It should be noted that neither of these challenges automatically refutes ontological reductionism, since there are various strategies that the reductionist could employ (such as invoking Boolean or mereological combinations to unify the 'many'), but they may shift the burden of proof on the reductionist or at least require that the reductionist position be further specified.

The sense of reductionism that I am interested in here is of a fairly weak type. I have labelled it 'ontological reduction', to distinguish it from 'logical', 'epistemic', or 'theory reduction', which in their various strengths suggest that the laws of biology follow deductively from the laws of physics or chemistry, or that knowing all the microstructural facts would enable one to know all the macrostructural facts. In contrast, ontological reductionism is a much weaker thesis, appropriately captured by Dupré ([2012], p. 131): 'dispositions of the whole are determined by properties of constituents, together with appropriate surrounding conditions' (see also Dupré [1995], Chapter 9; Gillett [2010]). Perhaps the thesis can be weakened even further, to the extent that ontological reductionism only implies a kind of supervenience thesis, whereby the biological is thought to supervene on the chemical (and the physical). This type of relation is a minimal requirement for providing a type of compositional explanation; Gillett ([2007], p. 195) calls it 'compositional reductionism', according to which one should commit to only 'one layer of truthmakers in the relevant component entities'. This could still be understood as a type of one-one reduction, where, say, a protein's three-dimensional structure is thought to supervene on an amino acid sequence. If we consider typical textbook models of protein structure, such as the structure of haemoglobin in

the biological function is and I will continue to use the notion of 'microstructure' for this purpose.

⁴ This should be distinguished from logical, methodological, epistemic, and theory reduction. Much of the anti-reductionist literature has focused on these stronger types of reduction, which may suggest, for example, that 'reducing one theory to another requires that we have some lawlike way to connect the vocabulary of the two theories' (Kincaid [1990], pp. 576–7). For discussion of different types of reduction regarding biology, see (Brigandt and Love [2017]), but notice that they focus on epistemic rather than ontological reduction. More on this below.

⁵ This is my own term for the phenomenon, developed in discussion with Jessica Wilson.

⁶ 'Realization' could be understood as an umbrella term that captures a variety of different dependence relations, including ontological reduction (see Gillett [2010]). We could also express this idea in terms of the 'nothing over and above' locution (compare with Wilson [2011]), that is, the biological function is 'nothing over and above' the chemical structure that realizes it. For a helpful survey of realization relations in metaphysics, see (Baysan [2015]).

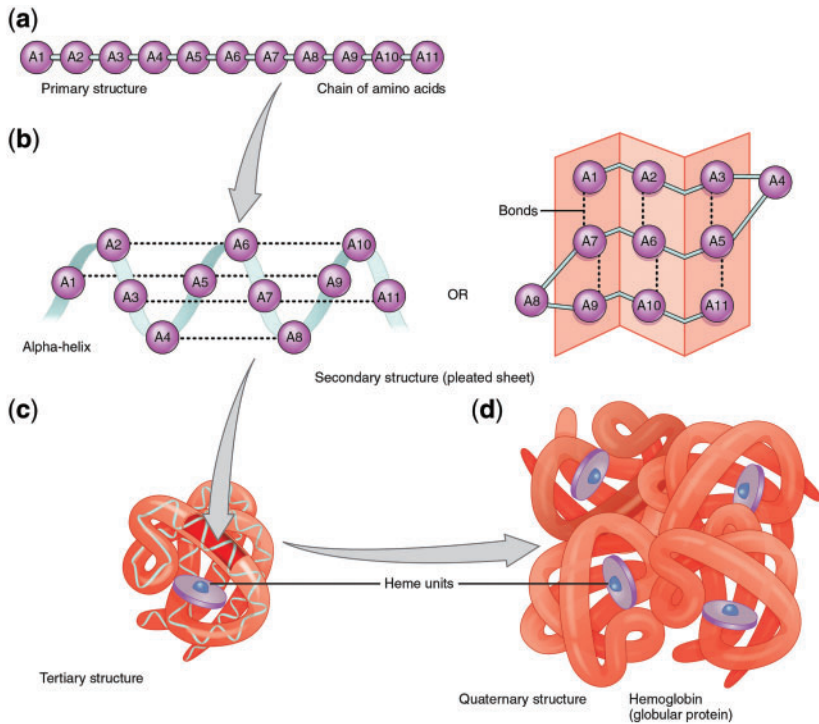


Figure 1. Textbook model of haemoglobin. Source: OpenStax College [2017]: ‘Organic Compounds Essential to Human Functioning’ (CC-BY 4.0), available at <legacy.cnx.org/content/m46008/1.8/>.

Figure 1, this type of idea is easy to understand. I will return to the structure of haemoglobin in Section 5, but we can already see that there are four levels to the structure of proteins and this structure may be straightforwardly described in terms of compositional reduction.

However, even this relatively weak form of reductionism may be resisted with an appeal to the phenomena of multiple determinations and multiple realization. Dupré himself points out that protein folding is a major problem even for this type of ontological reductionism. This is, perhaps, the consensus in the recent literature, although there are dissenters, such as Rosenberg (for example, [1985], [2006]). In what follows, I will examine the metaphysical underpinnings of the arguments from multiple determinations and multiple realization against ontological reductionism. These arguments do not, of course, automatically undermine the claim that the functional properties of proteins are determined by the microstructure. However, many authors (for example, Bartol [2016], pp. 542–2) continue to hold that multiple realization is at least a *prima facie* challenge for ontological reductionism precisely because

the functions are underdetermined by the microstructure. Before I discuss more specific examples, a few preliminary issues need to be addressed.

2 Functions: Aetiological or Goal-Directed?

First, it should be noted that a universal answer to the bio-chem question may not be available, or at any rate it will not be sought here. That is, even if some biological functions are reducible to microstructure, perhaps not all are. This is precisely why I will proceed by examining actual case studies. Second, it must be clarified what is meant by a 'biological function'. To avoid dwelling too long on the enormous literature on this topic, let me simplify the picture a little by contrasting two traditional positions: the 'aetiological' analysis and the 'dispositional' or 'goal-directed' analysis. The aetiological analysis emphasizes the historical path, the causal sequence that underlies the presence of a given function; it is 'backward-looking'. The goal-directed analysis emphasizes the role of a given function in the future, focusing on the biological fitness (reproductive success) of an organism resulting from the manifestation of a certain trait; it is 'forward-looking'. The aetiological analysis is the more popular of these two accounts, although the goal-directed analysis does have its defenders.⁷

One does not need to make an immediate choice between these approaches. This is because the choice between them may already reflect a given perspective. To see this, consider an observation that Mitchell ([1993], p. 259) makes in her critique of the goal-directed (dispositional) analysis:

The dispositional account has, if you like, weaker criteria or interest-dependent criteria for application. When the interests coincide with those of the evolutionary biologist, then function ascription will not differ greatly from the etiological theory. When the interests are different from those of the evolutionary biologist, however, then the two theories may deliver incompatible classifications of functional and nonfunctional consequences for a trait in question.

Mitchell distinguishes here between two explanatory goals, one of which reflects the aetiological analysis (the perspective of the evolutionary biologist), and the other reflects the goal-directed analysis. Note that this choice of perspective is in fact neutral about the ontological status of functions. Importantly, if we wish, we may understand functions as dispositional properties on either view, which is why 'goal-directed' may be a better label than 'dispositional' for the 'forward-looking' approach. To use a traditional example, the biological function of the heart is to pump blood and the heart has this disposition in virtue of its physical make-up, which is the result of an

⁷ A classic defence of the goal-directed, dispositional analysis is (Bigelow and Pargetter [1987]), but novel goal-directed accounts are emerging (see especially Austin [2015], [2017]).

evolutionary process. Depending on one's perspective, this account of the heart's function can be given either a backward-looking or a forward-looking analysis. Often the debate between these views concerns the relative merits of being able to explain why a given trait is present and being able to account for cases where a function does not manifest (see, for example, Bigelow and Pargetter [1987] and Millikan [1989] for discussion). Later (especially Section 6), we will see what kinds of considerations might speak in favour of one or the other approach.

3 Moonlighting and Multiple Determinations

My primary example will concern multiple realization, but let me first briefly consider the case of multiple determinations. In the recent literature, 'moonlighting proteins' have received special attention in this regard. These proteins perform multiple jobs or act in multiple roles; they are 'functionally multi-track' (Tobin [2010], p 51). This looks like a case of multiple determinations, one microstructure realizing multiple functional roles. Hence, we have a challenge for one-one reduction (see also Slater [2009]; Bartol [2016], pp. 536, 542). There are two types of moonlighting proteins: many change their functional role based on external, environmental factors, but some seem to lack a specific three-dimensional structure. Tobin ([2010], p. 42) labels these 'extrinsically structured moonlighting' and 'intrinsically unstructured moonlighting', respectively. There are reasons to think that only intrinsically unstructured proteins (IUPs) could be a serious threat to the ontological reductionist (Tobin [2010]; Dupré [2012], p. 134). The reductionist may insist that even though a protein's function may change depending on its environment, it is nevertheless the same microstructure that can enable these different functions. So, we may consider the various functional capacities of the protein to be just that: capacities—unactualized dispositions—that only manifest in the appropriate environment. This is fully compatible with the ontological reductionist's picture.

What about IUPs? The question concerns protein folding, and both Tobin and Dupré argue that the ontological reductionist is in some trouble in this regard. To simplify a little, it is the protein's folded, three-dimensional (tertiary) structure that enables its functional role. But proteins can be intrinsically unstructured in such a way that the same primary structure, the same sequence of amino acids, can fold up to produce different three-dimensional structures (compare with the simplified protein structure in Figure 1). This gives rise to what Tobin ([2010], p. 52) calls 'functional promiscuity': the same microstructure can result in distinct biochemical kinds. The question is, does the case of IUPs differ significantly from the type of moonlighting that was

deemed unproblematic for the ontological reductionist? Goodwin ([2011]) answers in the negative, so does Havstad ([2018]).

Goodwin's key insight draws from an analogy with organic chemistry, where we may also classify molecules into distinct functional types. Organic molecules can be classified depending on their molecular and reactive environments, but their behaviour in these environments is nevertheless considered to result from molecular structure (which should also be conceived of three-dimensionally). In other words, chemical kinds, for better or worse, are just as 'messy'—to use Havstad's term—as biochemical (or biological) kinds (see also Rosenberg [1985], p. 76). The case is analogous to the functional promiscuity of proteins. As Goodwin ([2011], p. 543) puts it: 'The tertiary structure of a protein, when there is one, represents the most energetically stable conformation available to the protein in the relevant biological circumstances'. The process of protein folding, which can produce different tertiary structures from one primary structure, may be understood as being governed by complex interaction with the environment, aiming towards the most energetically stable conformation. So, on the face of it, the ontological reductionist does have the necessary tools to deal with the moonlighting case. But is this enough to convince the anti-reductionist? Let me introduce some further metaphysical machinery to examine this case.

4 The Powers-Based Subset Strategy

Now that I have a case study involving multiple determinations at my disposal, I can briefly reconsider the ontological status of functions. Considering what we know of the phenomenon of protein folding, there is a potential benefit in interpreting biological functions as dispositional, causally relevant properties that refer to powers or potentials. By doing so, one can remain neutral about the goal-directed versus aetiological analysis of functions and also formulate a simple answer to the bio-chem question. Adopting this approach allows one to take advantage of an existing, versatile framework developed for the analysis of inter-level dependence relations: the 'powers-based subset strategy'. This framework was originally developed to define non-reductive physicalism in such a way that it can survive Kim's (for example, [1998]) challenge to non-reductive accounts.⁸ It also seems to fit the analysis of moonlighting proteins well.

So, we are trying to model the relationship between lower-level and higher-level phenomena in such a way that variation in the realization base and realized phenomena can be taken into account, not unlike in the context of

⁸ For a recent, elaborate account, see (Wilson [2011]; [unpublished]). Earlier versions can be found in (Wilson [1999]; Shoemaker [2001], [2007]); see also (Clapp [2001]; Shapiro [2008]).

non-reductive physicalism. The powers-based subset strategy is supposed to be neutral concerning the nature of the properties to be analysed, but a typical approach is to adopt some variant of the causal theory of properties, whereby the ‘powers’ may be individuated by the causal powers that individuals in which the properties are instantiated possess (but need not necessarily manifest) (Gillett [2010], p. 169).

In the present context, it may be tempting to adopt this framework because it would seem to be a good fit for the analysis of the potential of proteins to fold into various three-dimensional shapes, hence also accounting for the phenomenon of moonlighting. Following Goodwin’s suggestion, we could then associate these ‘powers’ with the protein’s primary structure. One reason for the initial appeal of the subset strategy for the ontological anti-reductionist is its potential to show that although higher-level properties are not identical to lower-level properties, they can still be ‘nothing over and above’ the lower-level properties. The way the subset strategy is supposed to show this is by demonstrating that the causal profiles of the higher-level properties are a subset of the causal profiles of the lower-level properties.⁹

Adapting the framework to the case at hand, one can formulate a potential answer to the bio-chem question as follows:

The set of causal powers associated with a biological function of a macromolecule are a proper (non-empty) subset of the causal powers of the microstructure (that is, its chemical properties) that realizes the function.

This answer to the bio-chem question is attractive, albeit simplified, because it is straightforwardly compatible with multiple realization and multiple determinations. For multiple realization, there could be several distinct proper subsets of the causal powers of the microstructure that can realize the set of causal powers associated with a given biological function. In the case of multiple determinations, there could be one proper subset of chemical properties capable of realizing two distinct functions, whereby the causal powers of the functional properties are a subset of a single chemical kind. Figure 2 illustrates this more generally.

The small black circles represent individual powers. On the left-hand side, we have two proper subsets (dashed circles) of these powers, with one overlapping power. If we imagine that we are tracking a function realized by that one overlapping power, then this will constitute a case of multiple realization.

⁹ I will sometimes simply talk about sets and subsets of properties as a shorthand for sets and subsets of the causal powers of properties. There are further relevant distinctions, for example, between ‘flat’ intra-level ‘property-realization’ and inter-level ‘microrealization’ (Gillett [2003]; Shoemaker [2007]). I am interested in inter-level realization between biological and chemical properties and will set these complications aside, following Kim’s ([2010], p. 105) simplified version of the subset strategy.

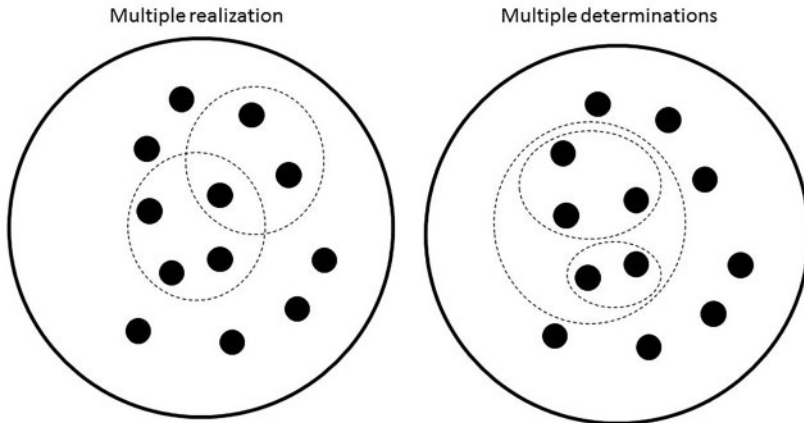


Figure 2. The powers-based subset strategy.

A scientific example of this could be a case where a very similar three-dimensional globular protein structure is produced by two distinct amino acid sequences. Even if the three-dimensional structures produced by the two sequences were not perfectly identical, they could have some of the same functional capacities (that is, powers). As we will see in the next section, the familiar haemoglobin molecule appears to be a good example of this type of multiple realization.

On the right-hand side, we are focusing on one larger subset of powers, with two further proper subsets (here depicted as non-overlapping, but there could be partial overlap as well). In this case, the larger subset of powers can realize two distinct functions, represented by the smaller subsets, hence constituting a case of multiple determinations. We may understand the larger subset as representing the collective causal powers of one biochemical kind. This idea can easily be applied to the case of moonlighting. Consider globular crystallin proteins (discussed in Tobin [2010], p. 42). These proteins have an important role in realizing a key power for the lens of the eye, namely, the refractive index of the lens (as described in Slingsby and Wistow [2014]). This role appears to have developed relatively late in evolution. Remarkably, many of these crystallins are identical to metabolic enzymes, hence serving a very different (and evolutionarily older) role, while also doing the job of a crystallin. The proteins are able to do these two ‘jobs’ by virtue of two distinct sets of causal powers, but both sets of powers are realized by just one biochemical kind, making this a clear example of multiple determinations. Slingsby and Wistow also provide a plausible evolutionary explanation for this phenomenon, suggesting that as proteins were recruited to serve as crystallins, it is likely that selection was made from among those proteins that already had some role in the cells of the

eye cup or surrounding tissues. So, it is not so surprising that these proteins may have retained their former roles in addition to serving as crystallins.

With these examples in mind, the powers-based subset strategy allows us to formulate the challenge for one–one reduction much more clearly, since it gives us a simple picture of multiple realization and multiple determinations. Applying this to the case of crystallins, the anti-reductionist could insist that even if the capacities of a crystallin protein are all included in one large set of lower-level powers—the protein’s primary structure—we cannot simply identify the higher-level powers of the three-dimensional macromolecule with this large set. Rather, it is only some proper subset (the smaller sets designated by dashed lines on the right-hand side of Figure 2) of the larger set that is relevant for each higher-level power, such as the two jobs of crystallins.

However, the ontological reductionist may legitimately ask: why should we focus on the proper subsets of only some of the causal powers of a biochemical (compare with Heil [2003]; Ney [2010]; Gillett [2010]; Wilson [2011], pp. 136–7)? In other words, are the various roles that a biochemical may act in just pragmatically selected from the overall causal profile of the biochemical kind? If so, then the advocate of ontological reduction could continue to argue that the further subsets of powers (as pictured on the right-hand side of Figure 2) are indeed reducible to the larger (dashed) subset, properly understood to include also those powers that may not always be manifested (depending on environmental factors). In this context, the larger (dashed) subset can be understood to include the set of powers that are contained in the molecular structure only as ‘capacities’. If this is right, then the powers-based subset strategy may be interpreted either according to the biological or the chemical perspective; it does not help to settle the debate between reductionism and anti-reductionism. Let me now move on to my primary example and add some further scientific detail.

5 The Case of Haemoglobin

My primary example is the haemoglobin molecule.¹⁰ Haemoglobin binds and transports oxygen in blood, releasing it in cells. Roughly speaking, the biological function of haemoglobin is its ability to bind and release oxygen; more accurately, to carry oxygen from the lungs to the tissues and return carbon dioxide from the tissues to the lungs. The haemoglobin molecule is constituted by four polypeptide chains: two alpha chains (alpha helices) and two beta chains (pleated sheets), which have different sequences of amino acid residues, but fold up to form similar three-dimensional structures (as in Figure 1). Only

¹⁰ This example is partly inspired by Kistler’s ([2016]; [2018]) recent work. Goodwin ([2011], p. 535) mentions the case in a footnote.

nine of the amino acid residues per globin chain in functionally identical haemoglobin molecules are present in all (functionally identical) haemoglobin molecules. Variations elsewhere in the sequence do not alter the functional profile. So, it appears that we are dealing with a case of multiple realization, where several different microstructures can produce the same biological function.

From the biological perspective, we can take two distinct macromolecules (polypeptides) as two instances of haemoglobin, whereby both macromolecules are effectively treated as the same protein, despite microstructural differences. Hence, 'haemoglobin' might serve as a label for a class of proteins defined not only in terms of their microstructural features but also in terms of their functional profile. But we need some more detail about protein structure to assess the example. It is at the level of tertiary protein structure, the three-dimensional combination of alpha and beta chains, where we see the biological function of the protein. There can be several different primary structures that produce the same functionality. This is a common feature of proteins. Significantly different primary structures can produce very similar three-dimensional structures, sufficiently similar in order to produce the same functionality. This is true in the case of the various haemoglobin molecules as well.

One aspect of variation in haemoglobin concerns its affinity for oxygen. For instance, since foetuses cannot breathe for themselves, the haemoglobin molecule active in foetal blood has a very high affinity for oxygen (there are three different haemoglobins in effect at different stages of foetal development). After birth, this affinity decreases. While it is important that haemoglobin has a high affinity for oxygen, it is also important that the oxygen gets released when needed. Accordingly, it is not surprising that the haemoglobins between adult mammals of different species may bear more similarity with each other than those between the adults and foetuses of the same species. In fact, haemoglobin can modify its affinity for oxygen depending on the chemical environment that it finds itself in. This behaviour is not unlike that observed with moonlighting proteins above, since the adaptability of haemoglobin is tied precisely to its environment. One property that influences haemoglobin's affinity to oxygen is the pH value of the environment.

I can now further specify the key bio-chem question. Are the capacities of the haemoglobin protein contained already in the primary structure? If they are present only at the level of tertiary structure, then we would appear to have a genuine case of multiple realization, which may be problematic for the ontological reductionist. But why should one think that various primary structures do not all have the same functional capacities?¹¹ Here we should recall Goodwin's point: the case of proteins is analogous to that of organic

¹¹ My thanks to an anonymous referee for suggesting this formulation of the question.

molecules that may be classified according to distinct functional types. Goodwin does not give detailed examples of this himself, but the classic case of acids (and bases), discussed, for example, by Stanford and Kitcher ([2000], pp. 115–9), makes for a good example.

Acids were historically defined in terms of their macroscopic, phenomenological properties, such as their sour taste and corrosiveness. But it proved difficult to come up with a proper microstructural characterization of acids. The most commonly used account derives from the Brønsted–Lowry theory: acid-base reactions generally involve the transfer of a proton (H^+) from an acid to a base. Hence, acids are proton donors and bases are proton acceptors. The obvious problem in giving a general definition for acids is that a variety of substances can undergo these reactions and indeed some substances (called ‘amphoteric substances’), such as water, can react both as acids and as bases. Quite clearly, this type of variation will depend on the environment. So, amphoteric substances are comparable to functionally promiscuous moonlighting proteins. Yet, it would be odd to claim that the capacity of an amphoteric substance to react as either an acid or a base is something over and above its microstructural properties, a capacity that the substance gains only when the relevant environmental circumstances are in place. After all, both losing a proton (reacting as an acid) and gaining a proton (reacting as a base) are simple chemical reactions that are determined by the relevant molecular structures and we understand these reactions very well. There is no reason to think that we could not accurately capture acid-base interactions in terms of molecular structure, as the capacity to act as an acid or a base is contained in that structure.

We arrive at the climax. Amino acids, which have a carboxylic acid group and an amino group (base) are also amphoteric. As we know, proteins are made up of amino acids and accordingly they can also react as amphoteric substances. So, it turns out that at least some of the interesting capacities of proteins, some of their functional promiscuities, derive precisely from their amphoteric nature, which we have just explained in terms of molecular structure. This looks like a straightforward case of ontological reduction.

6 Haemoglobin and the Problem of Lower-Level Vengeance

In this section, I will try to make it clear that there is more at stake in taking the microstructure rather than the functional profile of biochemical kinds to be the more fundamental classificatory criterion than it might at first seem. The reason is that there is a possibility that details relating to the microstructure will make a difference later on, even if they can safely be ignored for our current purposes. This can be called ‘the problem of lower-level vengeance’. The case of haemoglobin will once again serve as a good example.

In the previous section, we observed that there are nine shared amino acid residues per globin chain in functionally identical haemoglobin molecules. This is a relatively small number, as there are hundreds of amino acids in the four subunits of haemoglobin's primary structure. This 'remarkable similarity' between different haemoglobins has been previously noted¹²:

Although the hemoglobin molecules of all animal species differ in the order of most of the amino acids in their globin chains, they are identical (homologous) in structure at nine of the approximately 140 amino acids per globin chain. And it appears that just these nine amino acids principally determine the secondary and hence the tertiary and quaternary structure of hemoglobin as a whole. (Looijen [2000], p. 88)

If these nine shared amino acids indeed determine the three-dimensional structure of haemoglobin, then this gives us at least a *prima facie* reason to think that the causal powers of haemoglobin originate primarily in these nine amino acids. But perhaps even more importantly, each subunit of haemoglobin has an embedded haem group and each of these groups contains one iron atom. It is the haem group that has the ability to bind an oxygen molecule and hence the haem group is responsible for the most important biological function of haemoglobin. So, one might think that there is at least a one-one correlation between the functionally necessary basis of the nine shared amino acids and the haem group, and the biological function of haemoglobin. Importantly, the fact that the primary structures of haemoglobins may vary with regard to the remaining amino acids does not automatically make this a case of multiple realization.

Sometimes a change in a single amino acid is significant. In sickle cell anaemia, a single amino acid replacement (glutamic acid replaced by valine) in two of haemoglobin's four amino acid chains alters the functional profile of haemoglobin dramatically. This sensitivity of the protein's primary structure suggests that there may be important differences between the various haemoglobins even though they fulfil the functional role that we happen to be interested in (from the biological perspective), namely, binding and releasing oxygen. For instance, perhaps different primary structures could entail subtle differences in the stability of the relevant tertiary structure? This might not make a biologically (that is, evolutionarily) significant difference, but it will make a difference to the protein's chemical properties and hence its functional capacities. One such difference could be, say, the protein's heat sensitivity. Why? Because we know that proteins are very sensitive to heat, but slight differences in this heat sensitivity might not make any difference from the biological perspective (it could have no evolutionary significance).

¹² One can already find a more general account of the idea that the sequence of amino acids determines a protein's three-dimensional structure in (Rosenberg [1985], pp. 73ff.; see also Hüttemann and Love [2011]).

The protein may never be exposed to very high levels of heat, certainly not while it is still performing its normal function of carrying oxygen in the blood of living organisms. To appreciate this, we need to acknowledge more fine-grained distinctions in the class of haemoglobin molecules, and from the chemical perspective we can easily do so.

Here we reach the important point of this section: why should properties of the molecule that will not make a difference to the biological function be biologically interesting? Indeed, perhaps they should not, insofar as they truly do not make a biological difference, not even somewhere down the evolutionary line. Yet, it would require an extraordinary level of premonition to decide on this in advance, without researching what other differences there might be in the chemical properties of the various haemoglobin molecules. This is an issue where our understanding of functions will make a difference. Recall the distinction between the aetiological and the goal-directed analysis of functions from Section 2. For the proponent of the goal-directed approach, the potential future impact of even the slightest chemical variation could be significant. But for the aetiological analysis, only the causal sequence up to the present point is relevant. Either way, there is very good reason to be interested in such potential variation. Otherwise we might neglect differences that are clearly relevant even from the biological perspective. So, quite independently of the philosophical issue, this element of fallibilism could have direct empirical consequences that should not be masked by the choice between the goal-directed versus aetiological analysis of functions.

There is a philosophical consequence as well. If lower-level features such as chemical properties can turn out to have unanticipated consequences for higher-level phenomena such as biological processes, then there is a risk that the *prima facie* motivation to distinguish between these two levels diminishes. Often the motivation to postulate higher-level kinds as distinct from lower-level kinds is based on there being special science laws that refer to these kinds (for example, Polger and Shapiro [2016], p. 65; Wilson [unpublished], Chapter 3; we will return to this). But such special science laws are of course *ceteris paribus* laws, precisely because it can turn out that, say, subtle differences in the environment could cause exceptions. The question then becomes: what are the normality conditions in the case of the relevant biological laws? Whatever these conditions are, they should not be determined by the experimental setup alone. But this is precisely what is at issue in the case of haemoglobin. We should recognize that there could be differences in the chemical properties of different haemoglobins that could make a difference for biological processes as well. Otherwise, we may end up arbitrarily ruling out such biologically relevant differences based on nothing but the fact that we have not conducted any experiments where these differences would be observable. Therefore, an appeal to special science laws as the motivation to

postulate distinct higher-level kinds is not enough on its own. This is the problem of lower-level vengeance. Notice that this is also a general challenge for the powers-based subset strategy introduced in Section 4, because higher-level kinds may turn out to be dependent on lower-level kinds precisely in the sense that the ontological reductionist claims.

The upshot of my case study is this: classifying proteins like haemoglobin only in terms of their biological function is potentially too coarse-grained and involves a risk of interest relativity. This result reflects, to some extent, the work of Goodwin ([2011]), Bartol ([2016]), and Havstad ([2018]). Bartol ([2016], p. 543) recommends that we should adopt a ‘dual theory’ of chemical and biological kinds with respect to macromolecules such as proteins, driven precisely by issues surrounding multiple realization, whereas Goodwin’s view appears compatible with ontological reductionism. Bartol also points out that there could be at least pragmatic reasons to prefer a different classificatory practice. Further strengthening this line of thought, Havstad argues (partly against Goodwin’s reductionist line) that the primary structure of proteins, the sequence of amino acids, is not sufficient to account for protein individuation. She thinks that we need to consider the aetiological constraints as well and that these constraints are difficult to account for in microstructural terms, a worry that we have now seen repeatedly. Moreover, even if we knew everything there is to know about evolution and were able to reconstruct this in microstructural terms, we might nevertheless prefer a taxonomy of the relevant biochemical kinds driven by the biological perspective. The ontological reductionist may, however, still be able to accommodate this result, for the view is compatible with theoretical and semantic anti-reductionism, and the indispensability of higher-level predicates (as noted by Gillett [2010], p. 188). Even so, Havstad is right in arguing that the aetiological constraints are a challenge even for the most lenient microstructuralist. We will return to this challenge in Section 8.

To conclude this section, the answer to the initial question about the relationship between the biological function of a macromolecule and the microstructure that it is related to the bio-chem question will depend largely on whether one takes the biological or the chemical perspective in classifying macromolecules. Let me now try to generalize this result.

7 Multiple Realization and (Higher-Order) Interest Relativity

Multiple realization is widely considered to be a major hurdle for reductionism, not just in philosophy of mind, but in philosophy of science as well. We have now seen how this hurdle presents itself in the case of biochemical kinds and I have suggested a potential line of reply for the ontological reductionist. I will now consider what this implies for multiple realization more generally,

focusing on the recent account of multiple realization proposed by Polger and Shapiro ([2016]). Their latest account builds on earlier work and continues an interesting dialogue with Aizawa and Gillett (see, for example, Gillett [2003], [2007], [2010]; Aizawa and Gillett [2009]; Aizawa [2013]).¹³ Polger and Shapiro believe that there are fewer cases of genuine multiple realization than is often thought (though they do suggest there are some), arguing against the anti-reductive line. So, one might think that Polger and Shapiro's account makes for a natural ally to the ontological reductionist.

It is best to rely on Polger and Shapiro's own favourite example to illustrate some of their key ideas: the case of the multiple realization of a corkscrew (Polger and Shapiro [2016], pp. 24ff.; originally introduced in Shapiro [2000]). A corkscrew serves the function of opening a wine bottle by applying opposing forces to the lip of the bottle and the cork. In the fictional 'science of corkscrews', it is also stipulated that corkscrews must perform their function by means of a screw mechanism. Now, it seems clear that there are various ways in which the corkscrew 'kind' could be realized. The criteria seem to be fulfilled by a waiter-style corkscrew or a double lever corkscrew. The corkscrew kind is not supposed to be a purely functional kind, since some of the specific realizer properties are also specified, namely, the presence of a screw mechanism. So, not everything that is able to remove the cork from a wine bottle, such as a vacuum opener, will count as a corkscrew. This gives us a very simple case of multiple realization. Importantly, there may be other differences between the corkscrews that are not 'relevant' when it comes to multiple realization, such as the colour of the corkscrew or perhaps also the material that it is made of (as long as the device remains rigid enough). The central idea is that the same mechanism—the same function—can be performed by two or more systems in different ways (Polger and Shapiro [2016], p. 45). This is the starting point. My main interest is the enormous amount of work that the notion of 'relevance' does in Polger and Shapiro's theory.

It is important for Polger and Shapiro that not every difference contributes towards multiple realization. More generally, there are two ways in which a difference might be relevant¹⁴:

- (1) There may be a difference between two entities, A and B, that explains why A but not B can perform a certain function, that is, realize a certain property or kind.
- (2) There may be a difference between A and B that explains why, even though each is capable of performing a certain function, they each do so in a different way, making this a case of multiple realization.

¹³ Following Polger and Shapiro, my interest here is on the empirical thesis of multiple realization, not multiple realizability as a metaphysical possibility.

¹⁴ My thanks to Donnchadh O'Conaill for suggesting this distinction.

I shall focus on (2). In the case of the corkscrew, it was suggested that the colour of the corkscrew was not a relevant difference in the sense of (2), but the difference between a waiter-style corkscrew and a double lever corkscrew was considered a relevant difference. The question is: how do we decide which differences are relevant in this sense? This should remind the reader of the question posed at the end of Section 4, regarding the choice of the relevant proper subset. How can we distinguish lower-level kinds and higher-level kinds? The suggested answer is that the relevant differences are those that make for differences in kind according to the relevant special science. But if that is the case, then the corkscrew example is hardly going to take us very far towards a complete theory of multiple realization, as 'the science of corkscrews' is merely stipulated. That is why the real debate regarding multiple realization concerns genuine scientific case studies, such as the case of haemoglobin.

Which differences are relevant and why? Polger and Shapiro propose that often only relatively few differences are relevant for realization, whereas Aizawa and Gillett (for example, [2009]; Gillett [2003]) have argued that even very small differences can be relevant.¹⁵ To make any progress here, we must clarify one important issue. If the relevant differences are to be identified in terms of the relevant special science, then we must already assume that the special sciences are not interest relative, at least insofar as we are attempting to give an objective account of multiple realization. If multiple realization speaks to the issue of non-reductive physicalism and constitutes a challenge for ontological reductionism, an objective account is what we should be interested in. Yet, there is a risk that interest relativity of the type just described sneaks in. Polger and Shapiro ([2016], p. 65) are well aware of this risk:

[...] though this discussion of corkscrews invites the charge that our judgements of multiple realization depend on mere stipulations about which properties are important (e.g., rigidity) and which not (e.g., color), we believe this to be an artefact of the idealized example. In the case of legitimate scientific kinds, such as eyes, the sciences provide principled means for distinguishing relevant from irrelevant differences.

¹⁵ See also (Couch [2009a], [2009b]), where relevant and irrelevant differences concerning eyes are discussed; eyes are one of Polger and Shapiro's favourite examples and one of the few cases where they think that we have genuine multiple realization. Like Polger and Shapiro, Couch seems to think that genuine cases of multiple realization may be rare. Fang ([2018]) argues for the opposite conclusion and suggests that multiple realization in biology may be much more common once one adopts an understanding of realization along the lines of Gillett's 'dimensioned view', whereby the realized higher-level properties and the realizing lower-level properties may be instantiated in distinct objects and the causal powers of the higher-level properties may differ from the causal powers of the lower-level properties. Unfortunately, I do not have the space to discuss these proposals in detail here.

The basic idea behind this reply is certainly sensible.¹⁶ Genuine scientific cases are what matter, since it is the sciences that provide ‘principled means’ for distinguishing relevant from irrelevant differences; this is the desired source of objectivity in the account. But what are ‘legitimate scientific kinds’? Polger and Shapiro are allowed to refer to the sciences, but there is a risk that the problem of lower-level vengeance observed in the previous section will re-appear. Focusing on a higher-level science might lead us to disregard relevant lower-level differences. Even if one could somehow overcome this problem, there is a deeper problem with biochemical kinds. The problem is that, as has been observed, these kinds can be considered both from the chemical and the biological perspective.¹⁷ So, which science gets to decide what the relevant differences are? That is something that we just need to stipulate. We could label this the ‘higher-order problem of interest relativeness’. The fact that the special sciences can have conflicting ‘principled means’ and conflicting criteria for ‘legitimate scientific kinds’ produces different accounts for distinguishing relevant and irrelevant differences, resulting in higher-order stipulation.

We could take a step back here and consider which natural kinds science as a whole considers to be legitimate. But this would beg the question against the ontological reductionist who in fact wants to take precisely this view. The point is that from the ontological reductionist’s point of view, science as a whole ought to settle for a single account of a ‘legitimate scientific kind’, a type of natural kind monism, but it turns out that different special sciences may contradict each other in this regard. So, the ontological reductionist is entitled to attempt a reduction of higher-level kinds to lower-level kinds in the interest of reaching a pervasive theory of legitimate scientific kinds.

I will illustrate this key issue with a figure, once again employing the powers-based subset strategy introduced in Section 4. Figure 3 illustrates a simplified case of a biological function, which can be realized by two microstructures. The small circles represent powers and we can see that microstructures 1 and 2 share several powers (the overlap between the bigger circles), some of which are mutually sufficient to realize the biological function in question (the four powers at the overlap of the three bigger circles). From the biological perspective, it is the six powers overlapping the biological function that matter, they are the ‘relevant’ powers here. Each of the two microstructures has one power that is not shared, but which also contributes to the biological function. Applying this example to the case of haemoglobin, we could consider these additional powers to concern the affinity for oxygen (ignoring some scientific details for purposes of illustration), which is one of the powers that varies

¹⁶ It is an idea that has been repeated, in various forms, many times (see, for example, Fodor [1974], p. 102; Antony [2003], p. 9; Funkhouser [2007], p. 483).

¹⁷ For partly similar observations concerning the case of eyes, see (Couch [2009a], pp. 509–11).

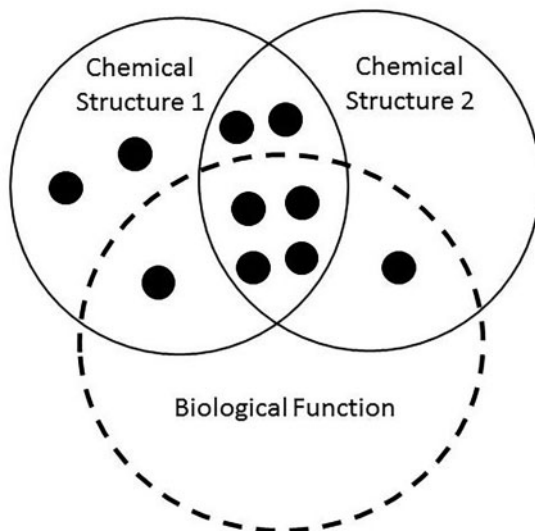


Figure 3. A simplified case of a biological function realized by two microstructures. The small circles represent powers.

between the various haemoglobins. So, it looks as if these two microstructures realize the biological function in relevantly different ways.

Is this a case of multiple realization or not? Well, from the biological perspective this presumably does constitute a case of multiple realization since the same function is realized by two different structures in relevantly different ways. The ‘relevance’ here concerns the important contribution that the difference between microstructures 1 and 2 makes (again, for a non-toy example, consider haemoglobin and the variation in the affinity for oxygen). But from the perspective of chemistry, there is little reason to think that this is a case of multiple realization, since there is, after all, a clear difference in the chemical properties (powers) that these two microstructures have. The fact that they are similar enough from the biological perspective to realize the same phenomenon is irrelevant. From the biological perspective, we have picked up on a shared subset of powers that the two structures have and that is indeed enough from the biological perspective, but it may not even occur to us to determine whether there are more fine-grained differences in the realization base. Yet, from the chemical perspective this case is entirely analogous to any overlap in the powers of two chemical kinds.

To drive this point home, consider any two arbitrary chemical compounds that happen to share a subset of their properties, such as their density, hardness, or something else. It is likely that these two compounds will share important similarities in their molecular structure as well, but just the fact that they share one or more of their chemical or physical properties does not

obviously entail anything about multiple realizability (in Polger and Shapiro's objective sense). After all, the shared subset of properties could be the result of a mere coincidence, a fluke resulting from the enormous number of different ways that elements can form compounds. However, for some (scientific) purpose, for some application or test that we may deem interesting, these two substances with a shared subset of properties could represent two relevantly different ways of doing the same thing, performing the same function, constituting a multiply realized, 'legitimate scientific kind'.¹⁸

The upshot of this account is simple. If the notion of relevance is relative to our interests, then we are not operating with an objective theory of multiple realization.

8 The Prospects for Ontological Reductionism

Let's take a step back and ask: why does multiple realization matter? Well, if biochemical kinds are multiply realized, then this constitutes at least a *prima facie* case against ontological reductionism. But why care about ontological reductionism? One motivation for pursuing ontological reductionism is that we might be able to reach a unified theory of natural kinds. This is surely one reason for the, perhaps fleeting, popularity of microstructuralism. And why strive for a unified theory of natural kinds? For some of us, the hope is that such a theory would also lay the basis for a unified theory of laws of nature and for scientific realism.¹⁹ The broader motivation has been aptly described by Gillett ([2016], pp. 34–40), who, citing the physicist and known defender of reductionism, Weinberg (for example, [1993]), emphasizes the sense of unity and parsimony that this type of ontological reductionism promises. This is not the place to pursue these grandiose goals, but perhaps it helps us to realign our focus.

With this focus in mind, consider where the discussion of biochemical kinds has led us. There is something of a consensus between Bartol and Havstad on this issue: the chemical perspective struggles to accommodate the evolutionary role of biological functions, the notorious aetiological constraints. This leads Bartol to conclude that we should postulate another group of kinds, evolutionary or historical kinds, following an established tradition in philosophy of biology. 'Are they natural kinds?', asks Bartol ([2016], p. 545). The answer given is: 'Like others, my argument for their naturalness will be left implicit, an appeal to the naturalness of natural selection' ([2016], p. 545). This is, of

¹⁸ Compare this result with Couch's ([2009b], p. 263) suggestion that 'we explain the notion of relevance in terms of the components of the structures that have to be appealed to in explaining the presence of the functions'.

¹⁹ For discussion, see, for instance, (Lowe [2006]; Sankey [2008]; Slater [2015]).

course, perfectly reasonable, but it is unlikely to convince the ontological reductionist. Relatedly, Havstad ([2018], pp. 20–1) concludes:

I'm not trying to say that a microstructuralist account of what confers membership in protein kinds is impossible to provide. Rather, I'm saying that doing so is at least going to be much more complicated than the microstructuralists have as yet admitted. It's going to have to be a disjunctive account, and somehow it's going to have to reconceive of at least one aetiological constraint microstructurally.

Havstad asks the proponents of microstructuralism to give a microstructural account of the aetiological constraint employed in distinguishing between proteins that share the same structure but may be distinguished due to aetiological constraints. Here the aetiological constraint seems to be a relevant difference in the sense discussed in Section 7. We would do well to accept this challenge, even though one may question whether the burden of proof is entirely on the microstructuralist (or ontological reductionist), given the earlier considerations. It is worth clarifying here that by 'a disjunctive account', Havstad means something comparable to her own and Hendry's ([2006], p. 872) disjunctive account of water, where the interaction of H₂O molecules in bodies of water is taken into account.

Rather than attempt to address Havstad's challenge directly, let me consider one more case study, a case upon which the ontological reductionist might hope to draw on in order to demonstrate why we might not necessarily want to reconceive aetiological constraints microstructurally even if we could: sickle cell anaemia.

It has been known for some time that heterozygous carriers of the sickle cell trait associated with sickle cell anaemia are protected from malaria. This is an obvious explanation for why that trait is overwhelmingly more common in sub-Saharan Africa—where malaria is also common—than elsewhere. Yet, this evolutionary story does not tell us anything about why the heterozygous carriers of the trait are in fact protected from malaria. The answer to the latter question was discovered only recently (Ferreira *et al.* [2011]). It turns out that the protection is due to the presence of carbon monoxide (CO) in the blood. The CO is a by-product of an enzyme that breaks down haem in blood. Because there are only low levels of free haem circulating in the blood of those with the sickle cell trait, the production of this enzyme is accelerated, resulting in increasing levels of CO (although not enough for toxicity). When malaria is added to the equation, the presence of CO limits haem levels in blood and this in turn appears to prevent the disease from developing.

So, there are two different explanations. The evolutionary story explains the prominence of the otherwise harmful sickle cell trait in sub-Saharan Africa, because it has a role in protecting its carriers from another harmful disease, malaria. However, the explanation for this causal power of the trait is not

itself aetiological, but ultimately something to do with the chemical properties of CO. Now that we know this, we can see that there is an intimate relationship between aetiology and the relevant chemical process. What has happened here is that the chemical process responsible for malaria protection was so elusive that we have only now started to comprehend it. Earlier, we had no choice but to merely gesture towards it in terms of the evolutionary story, but of course no one would deny that there must be some link between aetiology and these underlying chemical processes. Even though we now know what the relevant chemical process is, this does not mean that we should, or even could, replace aetiology-talk with chemistry-talk. Still, it is clear that we have made an important new discovery. We may now have the tools to develop new treatments for malaria. Indeed, Ferreira *et al.* ([2011]) discovered that it may be possible to treat malaria with CO.

In sum, Ferreira *et al.* managed to isolate the chemical process underlying the aetiological causal sequence associated with the current prominence of the sickle cell trait in sub-Saharan Africa. One of the chemical kinds possessing the relevant causal powers here turns out to be CO. I can now try to summarize the relevance of the sickle cell case. Before this recent discovery, there was no way to even get started with the microstructural explanation of the evolution of sickle cell anaemia. But now we have (at least most of) the elements that the microstructural explanation would require. In practice, we will probably never attempt to give this explanation, partly because it would just be a very long and boring conjunction of chemical processes. It would also be disjunctive, precisely in the sense that Havstad intends: the development of malaria can be terminated in other ways as well; CO is really just part of the story. Ferreira *et al.* ([2011], p. 406) also describe an alternative route that can prevent malaria from developing, which is based on sickle human haemoglobin exerting an immunoregulatory effect that inhibits cytotoxic T cell activation and expansion. Hence, a microstructuralist account that takes into account aetiological constraints would quickly grow inconveniently long and complicated.

Once we understand that a microstructural explanation is, in principle, available for the aetiology of sickle cell anaemia and malaria, no matter how inconvenient, there is no reason to think that similar explanations would not be available for all aetiological properties. If this is right, then there is no ontological difference between the biological and the chemical perspectives that could not be accounted for, even though we may never be in the position (epistemically or technologically) to actually give proper microstructural descriptions of many aetiological properties. One way to put this idea is to say that the difference between the two explanations is merely at the level of 'surface discourse' (it is primarily pragmatic) (Goodwin [2011], p. 535). Note that Goodwin does not appear to use this notion of surface discourse in

any pejorative sense. Rather, the idea is that the necessary preservation of the aetiological constraints at the level of surface discourse suggests a commitment to 'semantic anti-reductionism', whereby the predicates of the special sciences can still serve an invaluable, even indispensable role (see Gillett [2007], p. 195, [2016], p. 16). Nevertheless, differences at this level would not appear to be relevant, ontological differences in the objective sense required for genuine multiple realization.

There's no real surprise here. If aetiology can, in principle, be understood in terms of microstructure, then all that remains are epistemological, technological, and pragmatic considerations. These considerations make it clear that we will probably never attempt to give a microstructural explanation of most aetiological constraints, but the ontological reductionist only insists that it is possible to give such a story in principle. After all, we can easily introduce various abstraction levels of scientific language that serve some pragmatic purpose or other without any underlying ontological difference.

I can once more repeat the key questions that the opponent of ontological reductionism faces. How do we pick out the genuine higher-level kinds, how do we know what the 'relevant differences' are? One can attempt to rely on special science laws. The genuine properties and indeed genuine natural kinds are those that feature in special science laws (compare with Heil [2003], p. 209; Wilson [unpublished], Chapter 3). Heil rejects this picture, as I did earlier, and implies that on a liberal view of laws, such as Fodor's (for example, Fodor [1997]), the result is still undesirable. We end up 'reading off' features of the world from features of our language (Heil [2003], p. 218). I have suggested several reasons to think that this reaction is on the right lines. If there are biological laws distinct from the chemical laws that govern biochemical kinds, then the opponent of ontological reductionism may still have the upper hand. The most promising examples of this that I have discussed are the aetiological constraints governing biochemical kinds. But one ought to be careful not to confuse epistemic and metaphysical issues here. The aetiological story gives us a causal explanation of why the presence of a certain biochemical has been useful for the organisms that have it in their bodies, why it has helped them to survive and reproduce more effectively. However, there is a parallel story in terms of the chemical properties that underlie this causal sequence. We might not always know what this story is, but that is beside the point, since I am currently only interested in the prospects of ontological rather than epistemic reduction.

Let me be perfectly clear: Even if the ontological reductionist wins the day, none of this takes anything away from the practice of biology, nor does it undermine its status in any way. After all, the team (Ferreira *et al.*) that made the important discovery about sickle cell anaemia consisted primarily of biologists and health scientists. The status of a scientific discipline does not depend

on there being a set of genuine higher-level kinds or special science laws that we can associate with it. Aetiology and the historical chains that it is concerned with can be real even if the chains depend on microstructure. But the ontological reductionist insists that this microstructure is ontologically prior to aetiology.²⁰ This means that even if higher-order scientific language includes some indispensable classificatory practices, this does not entail pluralism about kinds (for example, Bartol's 'dual theory'), insofar as these different alternatives can be reconciled with the relevant microstructural, lower-level properties. We have seen that, at least in some cases, this can be done at the biology–chemistry interface. The upshot of this account is a type of ontological reductionism about biochemical kinds, which corroborates a unified theory of natural kinds instead of a dual theory of chemical and biological kinds.

Acknowledgements

I would like to thank Ken Aizawa, Umut Baysan, Carl Gillett, Max Kistler, Donnchadh O'Connell, Jessica Wilson, two especially helpful anonymous referees for this journal, and audiences at Auckland, Buffalo, Canberra, Otago, Prague, Sydney, Tampere, Toronto, Uppsala, Wollongong, and at the Third Society for the Metaphysics of Science Conference at Fordham University for their comments and suggestions. The research for this article was supported by two grants from the Academy of Finland (grant no. 266256 and no. 274715). Ongoing research on the article's topic is also supported by the European Research Council Project 'The Metaphysical Unity of Science' (grant no. 771509).

*Department of Philosophy
University of Bristol
Bristol, UK
tuomas.tahko@bristol.ac.uk*

References

- Antony, L. [2003]: 'Who's Afraid of Disjunctive Properties?', *Philosophical Issues*, **13**, pp. 1–21.
- Aizawa, K. and Gillett, C. [2009]: 'The (Multiple) Realization of Psychological and Other Properties in the Sciences', *Mind and Language*, **24**, pp. 181–208.

²⁰ Admittedly, as two anonymous referees point out, there remains some room for debate. It seems plausible that microstructure underlies historical chains in some sense or another. But one could attempt to argue that this is explanatory priority of a type that does not entail ontological priority, only some kind of pragmatic priority. I do not claim to have closed this possibility entirely, but I do believe that the account given above shifts the explanatory burden back on the opponent of ontological reductionism.

- Aizawa, K. [2013]: 'Multiple Realization by Compensatory Differences', *European Journal for Philosophy of Science*, **3**, pp 69–86.
- Austin, C. J. [2015]: 'The Dispositional Genome: *Primus Inter Pares*', *Biology and Philosophy*, **30**, pp. 227–46.
- Austin, C. J. [2017]: 'Aristotelian Essentialism: Essence in the Age of Evolution', *Synthese*, **194**, pp. 2539–56.
- Bartol, J. [2016]: 'Biochemical Kinds', *British Journal for the Philosophy of Science*, **67**, pp. 531–51.
- Baysan, U. [2015]: 'Realization Relations in Metaphysics', *Minds and Machines*, **25**, pp. 247–60.
- Bigelow, J. and Pargetter, R. [1987]: 'Functions', *The Journal of Philosophy*, **84**, pp. 181–96.
- Brigandt, I. and Love, A. [2017]: 'Reductionism in Biology', in E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, available at <plato.stanford.edu/archives/spr2017/entries/reduction-biology/>.
- Clapp, L. [2001]: 'Disjunctive Properties: Multiple Realizations', *Journal of Philosophy*, **98**, pp. 111–36.
- Couch, M. B., [2009a]: 'Multiple Realization in Comparative Perspective', *Biology and Philosophy*, **24**, pp. 505–19.
- Couch, M. B. [2009b]: 'Functional Explanation in Context', *Philosophy of Science*, **76**, pp. 253–69.
- Dupré, J. [1995]: *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Cambridge, MA: Harvard University Press.
- Dupré, J. [2012]: *Processes of Life: Essays in the Philosophy of Biology*, Oxford: Oxford University Press
- Fang, W. [2018]: 'The Case for Multiple Realization in Biology', *Biology and Philosophy*, **33**, pp. 1–24.
- Ferreira, A., Marguti, I., Bechmann, I., Jeney, V., Chora, A., Palha, N. R., Rebelo, S., Henri, A., Beuzard, Y., and Soares, M. P. [2011]: 'Sickle Hemoglobin Confers Tolerance to *Plasmodium* Infection', *Cell*, **145**, pp. 398–409.
- Fodor, J. [1974]: 'Special Sciences (Or: The Disunity of Science as a Working Hypothesis)', *Synthese*, **77**, pp. 77–115.
- Fodor, J. [1997]: 'Special Sciences: Still Autonomous after All These Years', *Philosophical Perspectives*, **11**, pp. 127–63.
- Funkhouser, E. [2007]: 'A Liberal Conception of Multiple Realizability', *Philosophical Studies*, **132**, pp. 467–94.
- Gillett, C. [2003]: 'The Metaphysics of Realization, Multiple Realizability, and the Special Sciences', *The Journal of Philosophy*, **100**, pp. 591–603.
- Gillett, C. [2007]: 'Understanding the New Reductionism: The Metaphysics of Science and Compositional Reduction', *The Journal of Philosophy*, **104**, pp. 193–216.
- Gillett, C. [2010]: 'Moving beyond the Subset Model of Realization: The Problem of Qualitative Distinctness in the Metaphysics of Science', *Synthese*, **177**, pp. 165–92.
- Gillett, C. [2016]: *Reduction and Emergence in Science and Philosophy*, Cambridge, Cambridge University Press.

- Goodwin, W. [2011]: 'Structure, Function, and Protein Taxonomy', *Biology and Philosophy*, **26**, pp. 533–45.
- Havstad, J. C. [2018]: 'Messy Chemical Kinds', *British Journal for the Philosophy of Science*, **69**, pp. 719–43.
- Heil, J. [2003]: 'Levels of Reality', *Ratio*, **16**, pp. 205–21.
- Hendry, R. [2006]: 'Elements, Compounds, and Other Chemical Kinds', *Philosophy of Science*, **73**, pp. 864–75.
- Hendry, R. [2012]: 'Chemical Substances and the Limits of Pluralism', *Foundations of Chemistry*, **14**, pp. 55–68.
- Hüttemann, A. and Love, A. C. [2011]: 'Aspects of Reductive Explanation in Biological Science: Intrinsicity, Fundamentality, and Temporality', *British Journal for the Philosophy of Science*, **62**, pp. 519–49.
- Kim, J. [1998]: *Mind in a Physical World*, Cambridge, MA: MIT Press.
- Kim, J. [2010]: 'Thoughts on Sydney Shoemaker's Physical Realization', *Philosophical Studies*, **148**, pp. 101–12.
- Kincaid, H. [1990]: 'Molecular Biology and the Unity of Science', *Philosophy of Science*, **57**, pp. 575–93.
- Kistler, M. [2016]: 'Espèces Naturelles, Profil Causal et Constitution Multiple', *Lato Sensu: revue de La Société de Philosophie des Sciences*, **3**, pp. 16–30.
- Kistler, M. [2018]: 'Natural Kinds, Causal Profile, and Multiple Constitution', *Metaphysica*, **19**, pp. 113–35.
- Looijen, R. C. [2000]: *Holism and Reductionism in Biology and Ecology: The Mutual Dependence of Higher and Lower Level Research Programmes*, Dordrecht: Springer.
- Lowe, E. J. [2006]: *The Four Category Ontology*, Oxford: Clarendon Press.
- Millikan, R. G. [1989]: 'In Defense of Proper Functions', *Philosophy of Science*, **56**, pp. 288–302.
- Mitchell, S. D., [1993]: 'Dispositions or Etiologies? A Comment on Bigelow and Pargetter', *The Journal of Philosophy*, **90**, pp. 249–59.
- Needham, P. [2011]: 'Microessentialism: What Is the Argument?', *Noûs*, **45**, pp. 1–21.
- Ney, A. [2010]: 'Convergence on the Problem of Mental Causation: Shoemaker's Strategy for (Nonreductive?) Physicalists', *Philosophical Issues*, **20**, pp. 438–45.
- Polger, T. W. and Shapiro, L. A. [2016]: *The Multiple Realization Book*, Oxford: Oxford University Press
- Rosenberg, A. [1985]: *The Structure of Biological Science*, Cambridge: Cambridge University Press.
- Rosenberg, A. [2006]: *Darwinian Reductionism: Or, How to Stop Worrying and Love Molecular Biology*, Chicago, IL: University of Chicago Press.
- Sankey, H. [2008]: *Scientific Realism and the Rationality of Science*, Aldershot: Ashgate.
- Shapiro, L. [2000]: 'Multiple Realizations', *The Journal of Philosophy*, **97**, pp. 635–54.
- Shapiro, L. [2008]: 'How to Test for Multiple Realization', *Philosophy of Science*, **75**, pp. 514–25.
- Shoemaker, S. [2001]: 'Realization and Mental Causation', in C. Gillett and B. Loewer (eds), *Physicalism and Its Discontents*, Cambridge: Cambridge University Press, pp. 74–98.
- Shoemaker, S. [2007]: *Physical Realization*, Oxford: Oxford University Press.

- Slater, M. H. [2009]: 'Macromolecular Pluralism', *Philosophy of Science*, **76**, pp. 851–63.
- Slater, M. H. [2015]: 'Natural Kindness', *British Journal for the Philosophy of Science*, **66**, pp. 375–411.
- Slingsby, C. and Wistow, G. J. [2014]: 'Functions of Crystallins in and out of Lens: Roles in Elongated and Post-mitotic Cells', *Progress in Biophysics and Molecular Biology*, **115**, pp. 52–67.
- Stanford, P. K. and Kitcher, P. [2000]: 'Refining the Causal Theory of Reference for Natural Kind Terms', *Philosophical Studies*, **97**, pp. 99–129.
- Tahko, T. E. [2015]: 'Natural Kind Essentialism Revisited', *Mind*, **124**, pp. 795–822.
- Tobin, E. [2010]: 'Microstructuralism and Macromolecules: The Case of Moonlighting Proteins', *Foundations of Chemistry*, **12**, pp. 41–54.
- Weinberg, S. [1993]: *Dreams of a Final Theory*, London: Vintage.
- Williams, N. [2011]: 'Putnam's Traditional Neo-essentialism', *The Philosophical Quarterly*, **61**, pp. 151–70.
- Wilson, J. [1999]: 'How Superduper Does a Physicalist Supervenience Need to Be?', *Philosophical Quarterly*, **50**, pp. 33–52.
- Wilson, J. [2011]: 'Non-reductive Realization and the Powers-Based Subset Strategy', *The Monist*, **94**, pp. 121–54.
- Wilson, J. [unpublished]: 'Metaphysical Emergence'.

