Diseases as Homeostatic Property Clusters

Chloé de Canson

Abstract. Several philosophers have recently drawn on property cluster accounts of natural kinds to argue that individual diseases form natural kinds. According to them, diseases have a super-explanatory property (their pathophysiology) in virtue of which their other properties (their symptomatology, biological signature, response to treatments, and prognosis) tend to co-occur. I argue that we can say more: disease pathophysiologies are *mechanisms*, such that the property clusters diseases form are *homeostatic*. I examine the two objections that natural kind theorists have raised against the homeostatic property cluster account, and I argue that they should not deter us from applying it to disease.

Understanding disease requires understanding not only what separates disease from health,[1] but also what separates individual diseases from each other. It seems clear that, although there may be significant differences between two instances of coronary heart disease, these two instances are more similar to each other than either of them is to an instance of multiple sclerosis. Moreover, these similarities and differences are not imposed onto the world by the human gaze: they are not mere postulates or conventions. Instead, they seem to already exist in the natural world, ready for us to discover them. This suggests that individual diseases form *natural kinds* (Fagerberg 2022, see also Dragulinescu 2010, Raburn 2024).[2] In this short paper, I argue that this idea is best developed by drawing on a particular theory of natural kinds, called *homeostatic property cluster* theory (Boyd 1989, 1991, 1999). To be clear, my aim here is rather modest: I do not argue that diseases form natural kinds, but merely that *if* one subscribes to the view that diseases form natural

[1] This topic has occupied philosophers of medicine for decades, with prominent views proposed by Boorse (1977), Wakefield (1992), and Cooper (2002). For an overview of this literature, see Sisti and Caplan (2016).

[2] Fagerberg (2023) also argues that the class of all diseases forms a natural kind. I do not take a stance on this separate question here.

kinds, one should subscribe to the stronger view that they form homeostatic property clusters.

Let us begin with Harriet Fagerberg's (2022) account of diseases as natural kinds. Fagerberg borrows from Ruth Millikan (1999, 2000, 2017) a "permissive" (Fagerberg 2022: 2) account of kinds. On this account, a kind is a category whose different instances share many properties, for non-accidental reasons. For instance, the category *Equus caballus* forms a kind because 1) its instances share a large number of properties—they have four legs and three gaits, they are herbivores, they have complex social lives, they are prey animals, etc., and 2) the correlation between these properties is not accidental, but results from the fact that horses are descended from other horses in a way that transmits these traits. There is no requirement that the instances of a kind share an exact set of properties. Indeed, there is presumably no fixed set of properties that it is necessary and sufficient to have in order to belong to the category *Equus caballus*. Some horses have fewer than four legs or more than three gaits but do not thereby cease to be horses. Instead, the properties that horses have are correlated with one another: they form 'clusters'. A 'property cluster' account of kinds such as Millikan's ensures that a degree of variability among the instances of a kind is possible without surrendering the kind's stability.

A kind's stability is secured by the 'non-accidental reasons' for which the kind's properties are clustered. Following Marion Godman et al. (2020), Fagerberg calls this the kind's *super-explanatory property*. For instance, the super-explanatory property of the kind *Equus caballus* is the property of being descended from another horse in a way that transmits various traits. Because instances of the kind *Equus caballus* have this property, they also tend to have a number of other properties (having four legs and three gaits, being herbivores, etc.) which Fagerberg calls *secondary properties*. The result is what Millikan calls a "clumpy world" (2017: 11): super-explanatory properties each give rise to clusters of secondary properties which are closely correlated among themselves, with large gaps between these clusters. There are plenty of animals with four legs and three gaits (*Equus caballus*), plenty of animals with two legs and five gaits (*Homo sapiens*), but no animals with two legs and three gaits. It is the fact that the world is clumpy this way which gives rise to our practices of classification; here, species taxonomy.

Fagerberg applies this account of natural kinds to diseases. She asserts that individual diseases are natural kinds: clusters of properties which correlate in virtue of a super-explanatory property. It is very important to note that this is a metaphysical and not an epistemological claim, which pertains to diseases as they actually are and not as we conceive of them. The claim is that diseases have a super-explanatory property in virtue of

which their secondary properties correlate, regardless of whether this super-explanatory property can be identified. According to Fagerberg, this super-explanatory property is the disease's "aetiology" (2022: 12). Ilana Raburn presents a similar account on which a disease's super-explanatory property is its "pathogenesis" (2024: 6). The disease's aetiology or pathogenesis produces a cluster of other properties: Fagerberg mentions the disease's symptomatology, and Raburn expands this to include the results of radiological or laboratory tests, the course of the disease, and the responses to treatment. So, on this picture, different instances of disease form a kind when they share an aetiology or pathogenesis, and in virtue of that, a presentation, a biological signature, a prognosis, and a treatment response.

As explained at the outset, my aim in this paper is not to defend the view that all diseases form natural kinds, but to argue that *if* they do, they have a further property—they form homeostatic property clusters. Nonetheless, the interest of my argument will be greater to the extent that the natural kinds view of diseases is rendered more plausible, and some conditions may be thought to constitute counterexamples to the view, for instance if they **(a)** have a poorly understood aetiology/pathogenesis, **(b)** have blurry boundaries, or **(c)** are defined by exclusion.[3] Let us therefore examine these potential counterexamples in turn.

> **(a)** Conditions with a poorly understood aetiology/pathogenesis constitute diseases on Fagerberg's view if there *exists* an aetiology/pathogenesis that explains their secondary properties. This is entirely consistent with this aetiology/pathogenesis being currently unknown or poorly understood. Thus, the mere fact that an aetiology/pathogenesis has not yet been found to explain the secondary properties of a condition does not entail that the condition fails to form a natural kind.
>
> **(b)** Some conditions, such as depression and anxiety, have considerable overlap. This raises the question whether they do indeed form separate natural kinds. Suppose (and this is contentious) that they do not, and that they thus fail to count as diseases on Fagerberg's view. It does not follow that psychiatric diseases do not exist. It merely follows that the real psychiatric diseases, whatever they are (and at least some psychiatric research is devoted to figuring this out), do not correspond to depression and anxiety as currently defined.
>
> **(c)** Nosological categories known as 'diagnoses by exclusion', such as fevers of unknown origin, group together instances of ill health that share a symptom and nothing further. These do not constitute natural kinds: by definition, the

[3] See e.g. Erasmus (2026) for a version of this worry applied to nosology.

> instances of these categories do not share secondary properties beyond symptoms, let alone an aetiology/pathogenesis. But they do not constitute a counterexample to the view in my opinion: the natural kind view of disease entails the highly plausible and commonly held judgment that such categories do not pick out bona fide diseases (e.g. Jutel 2010, 2011).

Note finally that meeting the criteria for a construct that does not or may not constitute a natural kind does not entail that one does not have a disease and is instead healthy. It simply entails that the disease one has does not or may not correspond to the construct that is currently employed to describe it. It would thus be entirely inappropriate to treat this person as if they did not have a disease.

This completes my exposition and modest defense of the view, drawn from Fagerberg, that diseases are natural kinds in a specific sense: each individual disease has a super-explanatory property that causes its secondary properties. This account of natural kinds, here applied to diseases, is known as the *simple causal view* after Carl F. Craver (2007) and Muhammad Ali Khalidi (2013). It states that the super-explanatory property of a kind *causes* its secondary properties to obtain, but does not specify which form of causation is involved. Now, the simple causal view can be strengthened, by specifying which form of causation connects a kind's super-explanatory property to its secondary properties.[4] For example, the *homeostatic property cluster* (HPC) account of kinds, formulated by Richard Boyd (1989, 1991, 1999), asserts that the form of causation involved is *mechanistic* causation: on the HPC account of natural kinds, the super-explanatory property of a kind is a *mechanism* that *underlies* the kind's secondary properties.[5] The view is called "homeostatic" to reflect the fact that the mechanism underlying a kind is similarity-generating: the secondary properties it produces form clusters. In the rest of this paper, I will illustrate the HPC account of natural kinds by applying it to disease, and I will argue that anyone who subscribes to Fagerberg's account of disease kinds, which is based on the simple causal view of kinds, can adopt a stronger view, one based on the HPC view of kinds. This stronger view amounts to the view that the super-explanatory property of

[4] This explains why Fagerberg characterised her account as "permissive" (2022: 2): it admits of possible strengthenings.

[5] Technically, the HPC view is somewhat broader, for it allows a second option: "Either the presence of some of the properties in [a family of properties] F tends (under appropriate conditions) to favor the presence of the others, or there are underlying mechanisms or processes which tend to maintain the presence of the properties in F, or both" (Boyd 1988: 323). Thus, all kinds whose super-explanatory properties are underlying mechanisms are homeostatic property clusters, but not all homeostatic property clusters have underlying mechanisms as a super-explanatory property. In this paper, I shall follow the rest of the literature and proceed as if the HPC account of kinds was narrower than Boyd made it to be, and stated that property clustering was always the result of an underlying mechanism.

diseases, what Fagerberg has called the disease's "aetiology" and what Raburn has called its "pathogenesis", constitutes the *mechanism* which underlies the disease's secondary properties: its symptoms, biological signature, treatment response, and prognosis. Furthermore, it amounts to the view that when two diseases are distinct, they are distinct because they are underlain by two distinct mechanisms.

Let us then attend to the nature of the disease's aetiology or pathogenesis. Raburn rightly points out that the term "aetiology" is sometimes used to refer to the historical cause of a disease (2024: 7).[6] For instance, smoking is sometimes said to feature among the possible aetiologies of lung cancer. But the super-explanatory property of a disease is not its historical but *underlying* cause.[7] Moreover, it is not the underlying cause of disease genesis only, but of the disease throughout its entire course. As a result, Raburn's proposal to replace the term "aetiology" by the term "pathogenesis" is also inadequate. Indeed, "pathogenesis" is sometimes used in a restricted way to only describe the process by which a disease initially develops. For instance, the pathogenesis of lung cancer is sometimes said to involve immune dysfunction that fails to keep early-stage malignancy in check. But the super-explanatory property of lung cancer is neither e.g. smoking nor e.g. early immune dysfunction. It is what *underlies* the secondary properties of the disease, not only in its early stages, but throughout its entire history. In other words, it is the uncontrolled division of cells of the lung. We are therefore looking for a term that, unlike "aetiology", refers to the underlying cause of a disease, and that, unlike "pathogenesis", refers to this underlying cause throughout the entire disease history. I suggest that we adopt terms often used in medical practice and in the medical sciences: *pathophysiology* and, less frequently, *pathomechanism*.

These terms, especially the latter, highlight something that neither "aetiology" nor "pathogenesis" made particularly salient, namely that the super-explanatory property for disease kinds is of a particular type: it is a *mechanism*. There are many ways to define a mechanism (see especially Glennan 1996, Machamer et al. 2000, and Woodward 2002; see Glennan 2017 for a minimal definition), but here, we will use a slight modification of Stuart Glennan's (2002) definition: for our purposes, the pathophysiology of a disease will be the complex system that produces this disease by the activity and interaction of its parts. On figure 1 below, the causal interactions between parts of the mechanism (P_i) and their

---

[6] Moreover, philosophers since Salmon (1984) have distinguished between *aetiological* (i.e. historical) and *constitutive* (i.e. underlying) causes.

[8] Smoking may cause lung cancer in one person and coronary heart disease in another: these have the same historical cause but different underlying causes (the uncontrolled division of cells in the lung for the former, and the reduction of blood flow through coronary arteries in the latter), and as a result, amount to different diseases.

activities (φ_i) are denoted with arrows, and taken together, these amount to a representation of the pathophysiology of the disease.

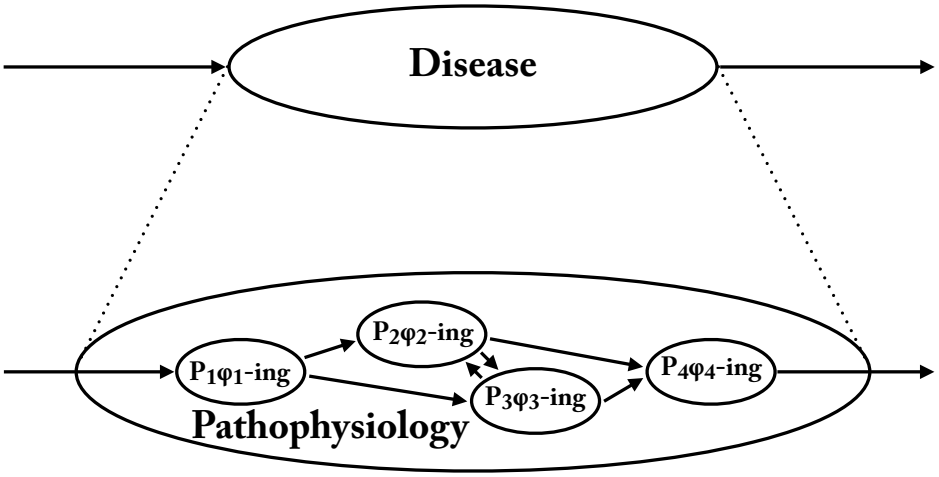

Figure 1. Representation of a disease's pathophysiology, adapted from Craver (2007).

Let us illustrate this with an example. The pathophysiology of multiple sclerosis, in very broad strokes, works as follows.[8] Immune cells known as T cells bind to and damage myelin, the sheath that insulates axons. (An axon is the part of a neurone that transmits signals from the cell body to the synapses, which then transmit these signals to other neurones.) These T cells also signal for further immune cells to travel to the site and damage myelin. As a result, large amounts of myelin are destroyed (this is known as *demyelination*), leaving behind plaques and impairing signalling within neurones. The pathophysiology of multiple sclerosis, as just sketched, explains its secondary properties:

- It explains why people with multiple sclerosis have the *symptoms* they do. For example, it is because plaque forms on the optic nerve that signalling is impaired therein and

[9] The illustration that follows is incomplete, in part because a perfectly detailed picture is not needed to illustrate my point, and in part because, as with any disease, there is much about its pathophysiology and its pathophysiology's relation to its secondary properties that remains unknown or poorly understood. One example: fatigue is the most debilitating symptom for many people with MS, but it is not currently understood how the underlying pathophysiology of MS leads to it. A second example: not everyone with any given presentation of MS reacts in the same way to particular treatments, but the reason behind these differential responses to treatments remain obscure. Now, I use aspects of MS that are well understood in my illustration, because they are simple to use, but as mentioned earlier in the paper, my claims apply to features of diseases both known and unknown. Thus, to say that MS is a disease on the natural kinds account of diseases is to say that the pathophysiology of MS (including aspects that we may not currently know) underlies all its secondary properties (including in ways that we do not currently understand).

that vision loss occurs. It is because plaque forms in the autonomous nervous system that signalling is impaired therein and that autonomic dysfunction, which can manifest as e.g. cardiovascular or bladder symptoms, occurs.

- It explains why people with multiple sclerosis have the *findings* they do. For example, it is because plaque forms that lesions appear on brain and/or spinal cord MRIs.
- It explains why people with multiple sclerosis have the *response to treatment* they do. For example, it is because the disease is driven by immune activity that many people with multiple sclerosis respond well to drugs that dampen it.
- It explains why people with multiple sclerosis have the *prognosis* they do. For example, it is because demyelination worsens with time (though, as we shall see, not in the same way for everyone), that people with multiple sclerosis experience increasing disability and, in some cases, death.

Thus, multiple sclerosis, *qua* disease, has several properties. It has the property of having a specific pathophysiology, that is, of being underlain by a specific pathomechanism. This property has a privileged status: it is super-explanatory, because it explains why all the other properties of the disease (its symptomatology, findings, treatment response, and prognosis) tend to co-occur.

Let us take stock. We followed Fagerberg in taking individual diseases to form natural kinds: they share a super-explanatory property, in virtue of which all their other properties correlate. This super-explanatory property, we have claimed, has a particular nature: it is a *mechanism*. It follows that it is an underlying mechanism that produces, and therefore explains, all the secondary properties of the disease kind: its symptomatology, findings, treatment response, and prognosis. Thus, if diseases are natural kinds, they are more than mere property clusters, they are *homeostatic* property clusters, property clusters whose super-explanatory property are similarity-generating *mechanisms*. We can therefore say more than Fagerberg had: diseases are natural kinds, not merely on the simple causal view, but also on the HPC view. This allows us to conclude that diseases are individuated by their underlying mechanisms: two instances of ill-health are instances of the same disease if they are underlain by the same mechanism (and thus have similar secondary properties), and they are instances of different diseases if they are underlain by different mechanisms (and thus do not have similar secondary properties).[9] However, there may exist reasons to resist drawing this conclusion. Two reasons are often given, in the natural kinds literature, for weakening the HPC account in favour of a more permissive view such as the simple causal view. Defending an HPC account of disease kinds will therefore require showing

[9] Thus, I defend what Erasmus (2026: §2.3) calls the "pathophysiological approach" to disease classification.

that neither of these reasons apply in the case of disease. In the remainder of this paper, I consider each of these two putative reasons against the HPC view in turn.

First, it has been pointed out that the super-explanatory properties of many natural kinds are not mechanisms (see e.g. Chakravartty 2007, Khalidi 2013, Slater 2015). For instance, a chemical element presumably forms a natural kind; yet, its super-explanatory property is not the fact that it has particular parts which are doing things (i.e. a mechanism) but the fact that it has particular parts *tout court*, namely, the fact that is has the specific number of protons captured by its atomic number. This raises the question whether the super-explanatory properties of *diseases* are necessarily mechanisms, or whether there may exist diseases whose super-explanatory properties are not mechanisms. Let us begin with two remarks.[10] First, the super-explanatory property of most paradigmatic cases of disease, including cancers, neurodegenerative diseases, cardiovascular diseases, autoimmune diseases, etc. is understood to be their respective known underlying mechanisms. Second, when what is known in medicine as a *symptom cluster* (a group of symptoms that tends to co-occur, have a similar prognosis, response to treatment, etc.; i.e. a cluster of secondary properties) is identified, biomedical science proceeds by attempting to identify the underlying mechanism that is assumed to give rise to this symptom cluster; thus, it is a regulative assumption of biomedical science that the super-explanatory properties of diseases are their underlying mechanisms.

Yet there appear to be potential counterexamples. Mechanisms *do* things. As such, mere states of affairs are not mechanisms. But several putative cases of disease seem to have states of affairs (and thus, non-mechanisms) as their super-explanatory property: explaining their secondary properties seems to require no more than pointing to a state of affairs. Let us consider two examples: the case of micronutrient deficiencies, and the case of injuries.

> *Micronutrient deficiencies.* The super-explanatory property of a micronutrient deficiency may seem to be not a mechanism but a state of affairs: the state of affairs whereby the micronutrient in question is deficient. Yet, as we shall see,

[10] It may also be interesting to note that Fagerberg's own *domino dysfunction* theory of disease (2025), according to which there is a disease where there is a dysfunction which causes other dysfunctions, seems to entail that diseases are always underlain by mechanisms: the mechanism whereby one dysfunction causes another which causes another. In the case of multiple sclerosis: an immune dysfunction (the inappropriate presence of T cells in the brain) causes another immune dysfunction (the inappropriate presence of other immune cells in the brain), which causes a neurological dysfunction (the destruction of myelin) which causes another neurological dysfunction (the formation of plaque on axons) which causes another neurological dysfunction (the impairment of neurological signalling), etc.

this thought dissolves when we distinguish micronutrient deficiencies per se, from the diseases associated with them.

Let us use iron deficiency to illustrate the point. It is entirely common for iron levels to be slightly low, and for there to be no symptoms at all. In this kind of case, the iron deficiency is inert: it is not associated with a mechanism, and it does not cause secondary properties besides a biological signature. The HPC account therefore rules it a non-disease. However, iron deficiency sometimes *is* associated with symptoms, and indeed, with a whole host of secondary properties. For instance, in iron deficiency anaemia, iron deficiency affects the production of haemoglobin, a protein in red blood cells implicated in oxygen transport. Thus, as iron levels decrease, the amount of haemoglobin decreases, and oxygen delivery to tissues becomes impaired. This pathomechanism explains the condition's secondary properties:

> *Symptoms*. It is because iron deficiency leads to a lack of oxygen in tissues that its symptoms (shortness of breath, dizziness, fatigue, etc.) arise.
>
> *Biological signature*. It is because iron deficiency leads to decreased haemoglobin levels that iron deficiency anaemia is characterised by low blood markers, not just of iron, but also of haemoglobin.
>
> *Treatment response*. It is because iron deficiency leads to low haemoglobin in red blood cells that blood transfusions, which supply new healthy red blood cells, can be a successful treatment.

Iron deficiency anaemia has a pathomechanism, and this pathomechanism explains its secondary properties: it is a disease on the HPC view.

The HPC account of disease therefore entails the following. Instances of *mere* micronutrient deficiency are not diseases, because they have no underlying mechanism, and no secondary properties (besides biological signature) for an underlying mechanism to explain. However, instances of micronutrient deficiencies can feature in mechanisms of diseases. For instance, iron deficiency features in the pathomechanism of iron deficiency anaemia. In such cases, it is not the micronutrient deficiency itself which constitutes a disease, but the micronutrient-associated condition whose pathomechanism does involve a micronutrient deficiency. Those who, like me, find this eminently plausible will conclude that micronutrient deficiencies do not constitute counterexamples to the HPC account of disease.

Injuries. The super-explanatory property of injuries (bone fractures, torn ligaments, wounds, etc.) appears to be not a mechanism but a state of affairs: the state of affairs whereby the bone is broken, the ligament torn, the skin open, etc. It seems indeed that in order to explain the symptoms or radiological findings associated with e.g. a bone fracture, one need only appeal to the bone's being broken. And the bone's being broken is not a mechanism. As a result, injuries do not count as diseases on the HPC account.

This may seem an issue for the account, since philosophers usually treat injuries as kinds of diseases: they use the term "disease" as an umbrella term for all that is pathological, in contrast to what is immoral, merely different, or a social problem (Cooper 2025). Although at odds with both ordinary and medical language, this usage can be explained by the fact that the primary goal of metaphysicians of disease has been to delineate the purview of medicine, that is, to draw a line between what medicine should concern itself with, and what it should not. The term "disease" has come to be used as placeholder for what *should* concern medicine, and the issue has come to be formulated as that of what separates disease from health.[11] This makes sense of their classification of injuries as diseases: since injuries *do* fall under the purview of medicine, they count as disease rather than health.

But the issue of what separates disease from health does not concern us here. Instead, what concerns us is the nature of the particular conditions which medicine considers diseases. It follows that we should use the term "disease", not as metaphysicians do when they study the purview of medicine, but as medical scientists and practitioners do when they participate in the investigation into and treatment of medical conditions. And since, on this use of the term, injuries do not count as diseases, they do not constitute counterexamples to the HPC view.[12]

Let us conclude on the first potential objection to the view that diseases are homeostatic property clusters: the objection that the super-explanatory properties of diseases may sometimes not be mechanisms. I have presented some general arguments against this objection, and defused two possible counter-examples to the HPC view. This of course

[11] See fn. 1.

[12] In fact, the HPC view may even explain the distinction between diseases and injuries that exists in medical language. The group of medical conditions (I am using this term as an umbrella term for anything that is of clinical relevance) which biomedical science investigates by attempting to understand their underlying mechanisms with the aim of intervening therein plays an important role in medical research and practice. It therefore makes sense for there to be a term that picks out just these conditions. According to the HPC account I am defending here, that is exactly what the the term "disease" does.

does not rule out the existence of further counter-arguments and counter-examples, but I hope to have made it seem plausible enough, that the onus to refute it is on the view's opponent.

This brings us to the second reason for which natural kind theorists have shied away from the HPC account. Boyd had proposed the HPC account as a middle way between essentialism (the view that kind membership is a function of having a particular necessary and sufficient set of essential properties) and conventionalism (the view that kind membership is a function of human conventions about how to group individual elements into batches) about kinds. His aim was to account for the fact that, despite a degree of variation among the members of a kind, their similarities are significant and exist independently of our practices of classification. In other words, Boyd had a *realist* aspiration: he wanted to establish that the (blurry) boundaries of natural kinds are assessment-independent. And he thought that could be achieved by indexing the boundaries of natural kinds to those of underlying mechanisms, such that if two property clusters are underlain by two different mechanisms, they constitute different kinds. But Craver (2009) then pointed out that the boundaries of mechanisms are not themselves assessment-independent. Depending on one's interest in describing a mechanism, one attends to different features thereof, in two ways: one includes certain parts and their activities in the mechanism and excludes others (the *boundary* problem), and one describes the mechanism at a particular level of abstraction (the *level* problem). I have termed these "problems" because, depending on the particular boundary and the particular level one adopts when describing a mechanism, the mechanism's individuation and typing varies. So, if the boundaries of natural kinds are to be indexed to those of their associated mechanisms, they will not have the assessment-independence that Boyd had initially hoped his account would secure.

Let us begin with the *boundary problem.* The pathophysiology of a disease, or its underlying mechanism, is the complex system that produces the disease by the activity and interaction of its parts. This raises the questions: Which elements, in the human organism, constitute *parts* of this mechanism, and which do not? Where does the complex system begin and where does it end? For example, the pathomechanism of multiple sclerosis as I have described it above begins with the binding of T cells on myelin, and ends with impaired signalling within synapses. May it not instead begin with what triggers this binding, and

end with e.g. bladder contractions?[13] The problem seems worrying, because different answers to these questions produce different answers to the question of what the disease *consists of*. Thankfully, mechanism theorists have discussed this issue at length. Craver et al., in an overview of the issue, write: "The philosophical puzzle is to articulate a principle by which entities, activities, and organizational features that contribute to the phenomenon are sorted from those that do not. The aim is to articulate a theory of *constitutive relevance*" (2025, emphasis added), and present three families of such theories.[14] Now, deciding on which theory of constitutive relevance to adopt is far beyond the scope of this paper, but by moving from a simple causal view to a homeostatic property cluster view of disease kinds we have nonetheless made progress. Indeed, we have been able to articulate the issue of what belongs to the super-explanatory property of disease in a way that Fagerberg and Raburn could not: this issue is the *boundary* or *constitutive relevance* for mechanism, an issue which already has a large literature on which it will be possible to draw.

This brings us to the *level problem*. Mechanisms can be described at different levels. For example, the description I gave above of the mechanism of multiple sclerosis is very high-level, and the mechanism so-described underlies all cases of multiple sclerosis. But there exist subtypes. For instance, in relapsing-remitting multiple sclerosis, particular immune cells known as regulatory T cells essentially repair part of the myelin that has been destroyed, in a process known as *remyelination*. But in primary-progressive multiple sclerosis, remyelination never takes place. This difference at the level of underlying mechanisms is associated with different secondary properties; for instance, people with relapsing-remitting multiple sclerosis have symptom exacerbation when demyelination occurs and partial symptom resolution when remyelination occurs, whereas people with primary-progressive multiple sclerosis never experience symptom reduction. Are relapsing-remitting and primary-progressive multiple sclerosis different diseases, or different subtypes of the same disease? The level problem laments the lack of assessment-independent answer to this question: if the mechanism is described at a higher level of abstraction, the two conditions count as a single disease, and if it is described at a lower level of abstraction, they count as two separate diseases. The question for us now is whether

---

[13] It is important to note that the boundary problem is distinct from what Peter Schwartz (2007) called the *line-drawing problem*. The line-drawing problem was first defined as the problem of what separates risk factors for disease from disease itself. It is now understood more generally as the problem of where to draw the line between sub-diseases, including so-called risk factors, and proper diseases (Binney 2025). The line-drawing problem is entirely orthogonal to the boundary problem: the former asks how severe a condition needs to get to count as a disease, the latter asks what physiological processes belong to a disease.

[14] The three families are the *mutual manipulability* account (Craver 2005, 2007), the *regularity* account (Couch 2011, 2023), and the *fat-handedness* account (Romero 2015). These are expanded on in Craver et al. 2025.

this constitutes a problem for the homeostatic property cluster view of diseases, a question which is especially pressing since one of the motivations behind the formulation of the view was its potential capacity, qua natural kind view, to individuate diseases.

Note that neither biomedical science nor medical practice assumes a completely assessment-independent account of disease. It is very common, in the biomedical study of one disease, to use another as a contrast class, and we see both multiple sclerosis contrasted to Sjögren's disease, and relapsing-remitting multiple sclerosis contrasted to primary progressive multiple sclerosis. It is also very common for medical practitioners to conceive of their patient as having multiple sclerosis *tout court* sometimes, and as having relapsing-remitting multiple sclerosis other times, depending on whether the context endows the difference to primary-progressive multiple sclerosis with practical import. According to biomedicine then, whether we consider the two types of multiple sclerosis discussed one mechanism or two mechanisms, one kind or two kinds, one disease or two diseases, will simply depend on one's practical considerations. Medicine has not made a decision either way: in some contexts, it will make sense to draw one line, and in others the other line. Now, one could object, of course, and insist that biomedicine errs in its individuation of disease, perhaps because one has a principled reason to maintain that disease kinds only exist at one specific level of abstraction. But in the absence of such a reason, it seems to me the natural thing to adopt the concept of disease operative in biomedicine. It follows that an account of disease kinds should entail that there is no assessment-independent fact about whether it is multiple sclerosis or its subtypes that count as diseases. The level problem is, after reflection, not a problem at all.[15]

Moreover, an account of natural kinds wherein levels are fixed is not necessary to produce an adequate *realist* account of disease kinds. Indeed, despite some interest-relativism in the choice between nested options, the HPC view imposes considerable constraints on what is allowed to count as a disease: it certainly does not collapse into a radically conventionalist account of natural kinds. At its most radical, conventionalism states that nothing in the world constrains how we draw lines around entities or properties. We can create a grouping of all the plants to which we are averse, call them *weeds*, and we will have created a kind. Similarly, on this view, we could generate a grouping composed of all instances of cardiovascular disease and of all instances of multiple sclerosis, and assert that we have a

[15] Note that the level problem arises at *all* levels of abstraction: we could lump multiple sclerosis with other T cell mediated autoimmune conditions, and we could split relapsing-remitting multiple sclerosis into molecular subtypes. Although rarer, one can imagine taking autoimmune disease, or MS-FCRL1 disease, to be individual diseases, when it is is one's interest to. For instance, we can imagine researching a treatment for autoimmune disease in general, and a treatment for just one molecular subtype of multiple sclerosis.

single disease. But the HPC view prohibits this. For according to the HPC view, all the instances of a kind must be underlain by a single mechanism. And while it may be assessment-relative whether one or two mechanisms undergird relapsing-remitting and primary-progressive multiple sclerosis, it is clearly not assessment-relative whether a single mechanism undergirds both cardiovascular disease and multiple sclerosis but no other condition. Such a mechanism does not exist. As a result, our gerrymandered disease candidate cannot count as a natural kind on the HPC view, nor as a disease on my view. This is exactly the kind of result we want an account of disease individuation to yield.

Let us conclude. We began with the idea that an individual disease forms a natural kind, whose super-explanatory property is its pathophysiology. Since the pathophysiology of a disease is a mechanism, we moved to the idea that individual diseases are homeostatic property clusters: their secondary properties tend to co-occur because their pathophysiologies are similarity-generating mechanisms. We then examined the reasons for which natural kind theorists have weakened the HPC view to the kind of permissive view we had started with, as applied to the specific issue of disease. Firstly, we investigated whether the super-explanatory properties of all diseases are always mechanisms, and we landed on the tentative conclusion that they do. Secondly, we discussed whether the bounded assessment-sensitivity of the HPC account counts against applying it to the case of disease, and we found that it does not. This leads us to the conclusion that we can say *more* than what theorists of disease kinds have said so far: it is not simply that diseases are property clusters, pace Fagerberg, but that they are *homeostatic* property clusters.

Acknowledgments

I am very grateful to Nick Makins, Camilla Trizio, and two anonymous reviewers for their particularly helpful comments on previous drafts, which allowed me to improve the paper considerably.

References

Binney, Nicholas. 2025. "The Line-Drawing Problem and Risk-Based Conditions." In *Handbook of the Philosophy of Medicine,* edited by Thomas Schramme and Mary Jean Walker. Springer.

Boorse, Christopher. 1977. "Health as a Theoretical Concept." *Philosophy of Science* 44(4): 541–573.

Boyd, Richard. 1989. "What Realism Implies and What It Does Not." *Dialectica* 43(1-2): 5–29.

———. 1991. "Realism, Anti-Foundationalism, and the Enthusiasm for Natural Kinds." *Philosophical Studies* 61(1): 127–148.

———. 1999. “Homeostasis, Species, and Higher Taxa.” In *Species: New Interdisciplinary Essays*, edited by R.A. Wilson, 141-186. MIT Press.

Chakravartty, Anjan. 2007. *A Metaphysics for Scientific Realism: Knowing the Unobservable*. Cambridge University Press.

Cooper, Rachel. 2002. “Disease.” *Studies in History and Philosophy of Biological and Biomedical Sciences* 33(2): 263–282.

———. 2025. “Fact, Value, and Disorder.” In *The Oxford Handbook of Philosophy of Medicine*, edited by Alex Broadbent. 591–613. Oxford University Press.

Couch, Mark B. 2011. “Mechanisms and Constitutive Relevance.” *Synthese* 183(3): 375–388.

———. 2023, “Clarifying the Relation Between Mechanistic Explanations and Reductionism.” *Frontiers in Psychology* 14: 984949.

Craver, Carl F. 2005. “Beyond Reduction: Mechanisms, Multifield Integration and the Unity of Neuroscience.” *Studies in History and Philosophy of Science Part C* 36(2): 373–395.

———. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press.

———. 2009. “Mechanisms and Natural Kinds.” *Philosophical Psychology* 22(5): 575–594.

Craver, Carl F., James Tabery, and Phyllis Illari. 2025. “Mechanisms in Science.” In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman. https://plato.stanford.edu/archives/win2025/entries/science-mechanisms/

Dragulinescu, Stefan. 2010. “Diseases as Natural Kinds.” *Theoretical Medicine and Bioethics* 31: 347–369.

Erasmus, Adrian. 2026. “A Goal-Directed Approach to Disease Classification.” *Studies in History and Philosophy of Science* 115(2): 102098.

Fagerberg, Harriet. 2022. “Reactive Natural Kinds and Varieties of Dependence.” *European Journal for Philosophy of Science* 12(4): 1–27.

———. 2023. “What We Argue About When We Argue About Disease.” *Philosophy of Medicine* 4(1): 1–20.

———. 2025. “A Domino Theory of Disease.” *Philosophy of Science* 92(2): 401–422.

Glennan, Stuart. 1996. “Mechanisms and the Nature of Causation.” *Erkenntnis* 44(1): 49–71.

———. 2002. “Rethinking Mechanistic Explanation.” *Philosophy of Science* 69: 342-353.

———. 2017. *The New Mechanical Philosophy*. Oxford University Press.

Godman, Marion, Antonella Mallozzi and David Papineau. 2020. “Essential Properties are Super-Explanatory: Taming Metaphysical Modality.” *Journal of the American Philosophical Association* 3: 1–19.

Jutel, Annemarie. 2010. “Medically Unexplained Symptoms and the Disease Label.” *Social Theory & Health* 8(3): 229–45.

———. 2011. *Putting a Name to It*. Johns Hopkins University Press.

Khalidi, Muhammad Ali. 2013. *Natural Categories and Human Kinds: Classification in the Natural and Social Sciences*. Cambridge University Press.

Machamer, Peter, Lindley Darden and Carl F. Craver. 2000. “Thinking About Mechanisms.” *Philosophy of Science* 67(1): 1–25.

Millikan, Ruth. 1999. “Historical Kinds and the ‘Special Sciences’.” *Philosophical Studies* 95: 45–65.

———. 2000. *On Clear and Confused Ideas*. Cambridge University Press.

———. 2017. *Beyond Concepts: Unicepts, Language, and Natural Information*. Oxford University Press.

Raburn, Ilana. 2024. “Intrinsic Kinds in Internal Medicine.” *Philosophy of Medicine* 5(1): 1–18.

Romero, Felipe. 2015. “Why There Isn’t Inter-Level Causation in Mechanisms.” *Synthese* 192(11): 3731–3755.

Salmon, Wesley. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.
Schwartz, Peter H. 2007. "Defining Dysfunction: Natural Selection, Design, and Drawing a Line." *Philosophy of Science* 74(3): 364–385.
Sisti, Dominic, and Arthur L. Caplan. 2016. "The Concept of Disease." In *The Routledge Companion to Philosophy of Medicine*, edited by Miriam Solomon, Jeremy R. Simon, and Harold Kincaid. Routledge.
Slater, Matthew. 2015. "Natural Kindness." *British Journal for the Philosophy of Science* 66(2): 375–411.
Wakefield, Jerome C. 1992. "The Concept of Mental Disorder: On the Boundary between Biological Facts and Social Values." *American Psychologist* 47(3): 373–88.
Woodward, James. 2002. "What Is a Mechanism? A Counterfactual Account." *Philosophy of Science* 69(3): 366–377.