Emerging into the Rainforest: Emergence and Special Science Ontology

Alexander Franklin* and Katie Robertson†‡

Forthcoming in the *European Journal for Philosophy of Science*

Abstract

Scientific realists don't standardly discriminate between, say, biology and fundamental physics when deciding whether the evidence and explanatory power warrant the inclusion of new entities in our ontology. As such, scientific realists are committed to a lush rainforest of special science kinds (Ross (2000)).

Viruses certainly inhabit this rainforest – their explanatory power is overwhelming – but viruses' properties can be explained from the bottom up: reductive explanations involving amino acids are generally available. However, reduction has often been taken to lead to a metaphysical downgrading, so how can viruses keep their place in the rainforest?

In this paper, we show how the inhabitants of the rainforest can be inoculated against the eliminative threat of reduction: by demonstrating that they are emergent. According to our account, emergence involves a screening off condition as well as novelty. We go on to demonstrate that this account of emergence, which is compatible with theoretical reducibility, satisfies common intuitions concerning what should and shouldn't count as real: viruses are emergent, as are trout and turkeys, but philosophically gerrymandered objects like trout-turkeys do not qualify.

Keywords: philosophy of science, philosophy of special sciences, autonomy, emergence, reduction, ontology, causal markov conditions

*alexander.r.franklin@kcl.ac.uk
†katie.robertson@stir.ac.uk
‡Authors listed in alphabetical order.

1 Introduction

If anything is real, viruses are real. That's because viruses enter into robust generalisations that are outstandingly empirically successful. And it is just this empirical success that motivates the scientific realist's commitment to them.[1] But viruses are not fundamental, and neither are cells, phonons or glaciers. Instead, the common expectation is that the fundamental constituents of reality will only be revealed by a yet-to-be-discovered final theory of physics. Of course, how to justify the scientific realist's commitment to non-fundamental entities and which brand of scientific realism is best is deeply controversial. Some of these controversies – e.g. those concerning the observable/unobservable distinction – are not our concern. What motivates us is that the heated debate about scientific realism *doesn't* hinge on whether scientific theories are or are not fundamental; instead this is the domain of the metaphysics of science. In this paper we present a new account of emergence, and thereby set out criteria that cast non-fundamental special science ontology as real. Our overall goal is that such criteria for admittance to the rainforest will satisfy the metaphysician, and yet will return an ontology appropriate to the scientific realist.

Yet maybe 'is real' should be applied to much more than just that which is dreamt of by all the sciences, even the non-fundamental ones, see e.g. Thomasson (2009). Some may want to say that in addition to every special science kind, every concatenation of spacetime points and every trout-turkey (the mereological fusion of a trout with a turkey) are also real. Here, the debate about the status of higher-level, or non-fundamental, ontology makes contact with van Inwagen's special composition question; the material plenitude view is committed to an explosion of entities (Fairchild (2019)), and so the entities of the higher-level sciences are bundled in alongside all possible mereological fusions.

Nonetheless there is an asymmetry between viruses (and special science kinds more generally), and those arbitrary concatenations; some discuss this metaphysical asymmetry, or privilege, in terms of naturalness (Dorr and Hawthorne (2013)), or sparseness of properties (Schaffer (2004)). Thus, another way to parse our project is as a discussion of which properties are natural, or of the 'scientific conception' of sparse properties – aiming to in-

[1] See e.g. Chakravartty (2017, §2.1).

clude 'being a virus' but exclude more gerrymandered properties.[2] Another idiom is that of entity realisation: we want to argue that while trout are realised in terms of the entities of scientific theories more fundamental than biology, we can be committed both to trout and their realisers without also committing to gerrymandered entities. Although we are content with these formulations of our view – and the reader should feel free to reformulate what follows accordingly – in the following we express our arguments in language closer to that of the scientific realist: we say that viruses are real but trout-turkeys are not.

A common way to cash out this metaphysical privilege is by appeal to theoretical irreducibility: while any properties of confected entities such as trout-turkeys are directly derivable from the bottom up, the suggestion would be that irreducible entities such as viruses have features that cannot be thus explained. However, given widespread theoretical reductionism many will be tempted to answer the question 'what is there?' with 'only the fundamental stuff'; see e.g. Barnes (2012). This is at odds with the criteria, methodologies, and sensibilities of the scientific realist. We set out and defend a view that resolves this tension by providing a principled basis for accepting the special scientist's ontology *and* granting that theoretical reductionism is rife.

One may analogise extravagant ontologies to rainforests on the one hand, and sparse ontologies to deserts on the other. Those attracted to desert landscapes stipulate that only the fundamental exists. But according to the scientific realist our reasons for accepting viruses into our ontology concern the quality of the evidence and explanations involving viruses, not whether viruses are fundamental or not. Indeed, were the scientific realist to follow the desert-lover's dictum of only committing to the fundamental, her present day ontological commitments would be so sparse as to be empty, since we lack a fundamental theory.[3]

But the desert is not to everyone's tastes and the gerrymandered kinds

[2] Schaffer (2004) argues that the fundamental conception, as opposed to the scientific conception, of sparse properties does well on the minimality criterion for sparse properties but that we should reject the minimality criterion; instead for Schaffer 'total science', i.e. the fundamental sciences and the special sciences together, give us the sparse properties. We agree, and see our project as similar in spirit though different in focus; our focus is on resolving the tension between scientific realism and ontological eliminativism.

[3] See Hoefer and Martí (2020) for an argument that the scientific realist should be committed to the entities of all sciences *other* than fundamental physics.

do not sit easily with the scientific realist. If ontological commitments are so easy to determine, why do scientific realists have to work so hard to establish the existence of particular scientific kinds?

The scientific realist wants the lush landscape of a rainforest but, like Goldilocks, she needs it to be 'just right': trout, viruses, phonons, and gases roam the rainforest, but trout-turkeys do not. The key issue – and the focus of this paper – is the admittance criteria for the rainforest. The terminology of 'rainforest realism' comes from Ross (2000), and although there are certain similarities between our accounts, discussed later, there are also significant divergences, the foremost being the role that reduction plays.

We say that even if viruses, phonons, and mountains can be understood entirely from the bottom up (that is, they are theoretically reducible), they are still in the rainforest just in case they are emergent. Provided that a non-fundamental or special science entity is emergent, it is inoculated against being eliminated by any theoretical reduction.

Here, our new account of emergence is compatible with reduction; this harmony has been suggested elsewhere by (e.g., Butterfield (2011a) and Crowther (2015)) but these are widely interpreted to be epistemological, rather than metaphysical, accounts of emergence (Barnes (2012) and De Haro (2019)). Our novel contribution is to draw out the metaphysical commitments of the scientific realist, and as such, our project is central to naturalised metaphysics.

In section 2 we argue that irreducibility (understood as the failure of inter-theoretic reduction) is the wrong peg on which to hang the existence of the higher-level entities. Section 3 outlines our account of emergence: for us, emergence requires screening off as well as novelty. Section 4 applies our account to the case study of viruses, and section 5 demonstrates that trout-turkeys and other gerrymandered examples fall foul of the criteria and, thus, are denied admittance to the rainforest. In section 6, we relate our account to prior analyses of rainforest realism, and show that our view encompasses that of Ladyman and Ross (2007), but avoids some of the criticism levelled against their project. In section 7, we conclude.

2 A Traditional View

'Irreducibility' is often taken to be key to the existence of special science entities, for example, Ney (2021) says: "for the reductionist, the hope is that for all phenomena, they will either be identified with entities of physical science or eliminated altogether in favor of the entities of a superior theory".[4] Indeed, irreducibility is very closely related to the name of one position committed to the existence of the special science entities: non-reductive physicalism. According to this view, the higher level (or special science) supervenes on the fundamental physical level, but is not reducible to it.[5]

However, there are at least two distinct ways to understand what 'irreducibility' means. First, irreducibility is a relation between different theories, see e.g. Butterfield (2014) and Nagel (1961). Second, irreducibility may be read ontologically; thus, the irreducible entities are those which exist *in addition* to the fundamental entities. On this ontological conception, the reduction of the mental to the physical, or biological to the chemical *is* the claim that there are no mental (or biological) entities in addition to the physical (or chemical) entities. But this doesn't offer any explanation for the contents of our ontology. So, while we accept that ontologically irreducible entities will be admitted into the rainforest, henceforth we'll understand (ir)reducibility as a claim about inter-theoretic relations.

The crucial point is that, on our view, theoretical reducibility is divorced from the criteria for realism/admittance to the rainforest, thus our project should be of central interest to those who are realists about higher-level entities even if they hold that these are theoretically reducible.

On Nagel's influential account of inter-theoretic reduction, one theory T_t is reduced to another T_b if the descriptions given by T_t can be deduced from T_b, perhaps with the aid of some additional assumptions (known as bridge laws).[6] This deduction allows the entities and their behaviour described by

[4] See also e.g. Gillett (2016).

[5] The connection between the ontological status of the special science entities (such as economies, magnets and viruses) and non-reductive physicalism is controversial. Some, such as Barnes (2012), claim the two topics are orthogonal. Others, such as Crane (2001), claim that these topics are the same debate under a different name.

[6] For other related accounts see e.g. Dizadji-Bahmani, Frigg, and Hartmann (2010), Rosaler (2017), and Schaffner (2013).

T_t to be explained from the bottom up.[7] Note that this is not to suggest that the higher-level reduced theory is, thus, redundant or could be eliminated; only that its adequacy to its target phenomena can be explained.

Given this characterisation of reduction, we can now state the traditional view that we oppose: the view that it's only (theoretically) irreducible entities that exist. Nagel (1961, p. 337) notes this view: "successful [inter-theoretic] reduction of thermodynamics to statistical mechanics was [...] taken to prove that spatial displacements are the only form of intelligible change, or that the diverse qualities of things and which men encounter in their daily lives are not 'ultimate' traits of the world and are perhaps not even 'real'".

Our view is that non-fundamental entities exist if they are ontologically emergent, and much of this paper is devoted to the detailed analysis of what it takes to count as ontologically emergent, but we resist the conflation of ontological emergence with theoretical irreducibility for the following reasons:

First, (ir)reducibility, as stated, is a relation between theories. While this is a guide to the relations between the entities, there's no completely straightforward relation between theories, understood as scientific representations of the world, and the entities that we find in the world. On our view, theories describe the phenomena, and entities participate in such phenomena, where 'entity' is neutral between objects, states, and events (see e.g. Bogen and Woodward (1988, p. 321)). As such, it's not easy to make theoretical irreducibility serviceable to distinguish between ontologically emergent and ontologically eliminable entities.

Second, while the theoretical reducibility of entities can be called into question, the supervenience of special science entities is far less controversial.[8] However, in certain circumstances reduction and supervenience are not far from one another. Indeed, when the conditions of Beth's theorem hold, supervenience collapses into reduction, and so maintaining supervenience *without*

[7]Nagel's account is closely related to, and inspired by, Hempel's deductive-nomological view of scientific explanation, so for Nagel reduction and explanation are interwoven. On other accounts of scientific explanation, a deduction of T_t from T_b will not be so closely linked to explanation.

[8]Even those who advocate strong emergence, such as Hendry (2010), aim to maintain the supervenience of the special sciences such as chemistry on the physical; though see Seifert (2020) for the suggestion that Hendry's position is unstable.

reduction is an unstable position.[9] Of course, the assumptions required for Beth's theorem do not hold for many realistic cases of scientific theories, since our scientific theories are not first-order formal languages. Should we discover that we can formulate a pair of scientific theories in a first-order language, then it seems that the assumption of supervenience would entail reduction, and so vanquish the commitment to the higher-level entities. But we do not expect our commitment to the existence of viruses, mountains, and magnets to be sensitive to such formal developments.

Third, even if viruses and suchlike can be reduced (since the theory describing them can be understood from the bottom up), we don't think they should be ejected from the ontology. So where we have inter-theoretic reduction, we needn't then follow through with ontological reduction, or elimination of the higher level. Widespread reduction looks like a plausible bet to us.[10] Later we will see that the case study of viruses provides an exemplar of a kind (or many kinds) that certainly should be taken to exist, but may well be reducible. Moreover, even if viruses can't be understood entirely from the bottom up, that's not the reason that they exist.

Our account of ontological emergence is compatible with theoretical reducibility. Ontological emergence then licences admission to the rainforest and is the ground of ontological irreducibility.

While we claim that ontological emergence, rather than theoretical irreducibility, is the admittance criterion to the rainforest, reduction is still relevant for classifying types of emergence: emergence that's compatible with reduction is termed 'weak emergence', whereas emergence combined with the failure of reduction is termed 'strong emergence'.

Although we believe that inter-theoretic reduction is plausibly widespread, it is not a requirement of our account. As such, those (like Bedau and Humphreys (2008) and Chalmers (2006), and us in this paper) who

[9]The first philosophers to emphasise Beth's theorem as threatening such a collapse were Hellman and Thompson (1976); Butterfield (2011a) discusses this topic in detail.

[10]In keeping with the original work on reduction since Oppenheim and Putnam (1958), we agree that whether reduction holds in particular cases is an empirical matter, and the plausibility in any given case depends on how stringent one's detailed account of reduction is. Nonetheless, we see the success of research programs such as the human genome project, work in biophysics and many other areas as reasons to be optimistic that reduction will succeed more generally, see e.g. Schaffner (2013) and Waters (2008).

define strong emergence in terms of failure of inter-theoretic reduction, will take our focus to be primarily, though not exclusively, on weak emergence. On the other hand, those like Gillett (2016), who take the strong/weak emergence distinction to correspond, respectively, to ontological/epistemological emergence will take our focus to be on strong emergence. Either way, our examples of ontological emergence are, in principle, compatible with reduction, but our account should still be relevant to those interested in (putative) failures of reduction.

We began this section with Ney's contention that the reductionist aims to choose between elimination and identification. In what follows, we focus our critical attention on the eliminative branch, but note that many – especially those following in the tradition of David Lewis – plump for identification. To our eyes, the practice of science furnishes us with very few examples of straightforward identification of entities and properties; see e.g. M. Wilson (1985) for a critique of the temperature-kinetic energy putative identity.[11] Ultimately, we are sceptical that that all the scientific realists' ontological commitments can be understood to be identical with the ontology of fundamental physics, so we think it worthwhile to explore an alternative view.

3 Our Account of Emergence

> **Definition**: An entity is emergent if and only if it is involved in dependencies that are novel and screen off lower-level details.

According to us, emergent entities feature in macrodependencies that have two key features. Firstly, the macrodependencies must screen off the microdetails in a certain sense, to be outlined in section 3.1. Secondly, the macrodependencies must be suitably novel, as discussed in section 3.2.

But what are the macrodependencies that we have in mind? Macrodependencies are projectible; each macrodependency may fall into various cat-

[11] Another naturalistic worry for identification is that dynamical processes, interactions, and forces are essential to the formation of macroscopic entities. The mereological composition of certain parts never suffices for the whole, and interactions and processes occur over time, so the identification would have to be over temporally extended segments; see Ladyman and Ross (2007).

egories, but it can be useful to distinguish the following. They can be dynamical: as is the case with the differential equations for reaction rates in chemistry, population dynamics in biology, or the equations describing the approach to equilibrium in statistical mechanics. There are also causal macrodependencies: perception of predators by prey causes flight, or a change in temperature causes a change in magnetisation. And often the macrodependencies are law-like: for example, the ideal gas law is an example of a nomic macrodependency, as are the laws of the special sciences such as economics' law of supply and demand. By 'macrodependency' we do not place any constraint on the scale at which the dependency operates, only that it is not a dependency featuring in any putative fundamental physical theory.

Note that, despite discussing scientific representations or descriptions, we aim for our account of emergence to be one of ontological emergence. Accounts of emergence in the philosophy of science are often dismissed as merely of epistemological interest (cf. Barnes (2012), as discussed by Crowther (2015)), and clearly accounts that define (weak) emergence as involving 'surprise' or 'unexpected features' render emergence a feature of our knowledge of the world (Chalmers (2006)). But we are interested in emergence as a feature of the world rather than our descriptions of the world. That being said, the best way we can get a handle on such worldly features is through our scientific representations of the world – scientific theories. Consequently, we refer to the 'microdetails' that correspond to features both in particular scientific descriptions and in the facts they represent.

In addition, we assume that the macrodependencies in question are genuine macrodependencies. That is, in keeping with the attitude of the standard scientific realist, we presume that there are higher-level laws, higher-level causal relations, and higher-level dynamics.[12]

In general, emergence balances some form of dependence on the more fundamental with some form of independence, autonomy or distinctness.[13] Our account of emergence is slightly unorthodox in that we won't have one

[12]The assumption that there are genuine special science laws, dynamics, and causes is part and parcel of the scientific realist's view, but note: this presumes that a satisfactory solution to Kim's causal exclusion problem has been found, e.g. by explaining why overdetermination is non-problematic (cf. Bennett (2007)).

[13]Taylor (2015, p. 654): "emergent properties are macro-level properties that are in some sense both dependent on and autonomous from their underlying micro-level properties".

criterion for the form of dependence and a separate one for the form of independence. Both will be present in the first of our two conditions for emergence: screening off. Macrodependencies that screen off microdetails are important because they are central to understanding how higher-level facts can be autonomous (in a particular sense) from the underlying microfacts.

3.1 Screening Off

Entities count as emergent in our sense if they enter into macrodependencies that have two features: novelty and screening off. The macrodependency in question must screen off some lower-level details that might seem – at first glance – highly relevant to the phenomena in question. This screening off condition has two parts: a dependence part – 'unconditional relevance', and an independence part – 'conditional irrelevance', following Spirtes, Glymour, and Scheines (2000) and Woodward (2021), as developed by Robertson (n.d.).[14]

Let's start with an example of a nomic macrodependency: *under usual background conditions, water boils at* $100°C$. This is a macrodependency B-A: where B is the temperature, and A is the boiling of the kettle. We will use this to illustrate the key idea of screening off, and the connection to a particular account of autonomy from which we take the core ideas.

First, let's consider the dependence part. While there's a macrodependency B-A describing the goings on in the kettle, there are other relevant details – in particular there are lower-level details that B and A supervene upon. In this example, water boiling depends on the underlying properties: the molecular motion of the water molecules is relevant to whether the kettle boils (if they are moving faster, then the kettle is more likely to boil). This is the dependence part of the condition.

We can spell this out following the framework discussed by Robertson (n.d.) and Woodward (2021) in terms of 'unconditional relevance'. The lower-level description, or microstate, of the water is relevant to whether the kettle boils. There is a probabilistic fingerprint of this dependence but to make this precise we need to specify what the microstate or lower-level description (LLD) is. For now, we take the LLD to be the supervenience basis of the

[14]See Sober (1999) for an appeal to screening off with goals related to ours.

temperature (B). (We will come to the supervenience basis of A shortly).

Since (in)dependencies can be represented using probabilities, we can define unconditional relevance as follows.

> **Unconditional relevance:** conditional on a particular lower-level description, the probability of the macro-description A obtaining increases: $P(A|LLD) > P(A)$. Under certain circumstances,[15] $P(A|LLD) = 1$.

But lots of the lower-level fine-grained details don't matter, i.e. they don't influence the macrodescription. One way of putting this: those details can be changed without influencing A. So there are other descriptions LLD_1, LLD_2,...,LLD_N for which $P(A|LLD_i) > P(A)$. In the kettle example, there are a range of distinct molecular motions that give rise to the kettle boiling – exactly how the boiling is realised doesn't matter. (E.g. swapping molecule 105830's position with molecule 300594's position makes no difference). We can think of these different lower-level descriptions as forming an equivalence class, a collection of lower-level descriptions that subvene B and influence macrofact A.

On to the independence part: for now, let's assume that which member of the equivalence class of microstates is instantiated isn't going to matter to the occurrence of A (we will relax this assumption shortly). This foreshadows how the lower-level details are irrelevant – roughly, given the macrostate B (namely $100°C$), the molecular motion is irrelevant.

As such, molecular motion is unconditionally relevant to the boiling of your kettle – that is, the boiling depends on the molecular motion. But once we conditionalise on the temperature of the water, any further details of the molecular motion are irrelevant. This is known as conditional irrelevance. While we are operating under the assumption that B-A is a strict macrodependency, this is then a case of exact screening off and we can define conditional irrelevance:

> **Conditional irrelevance (strict version):** $P(A|B\&LLD) =$

[15] If the microdynamics take all the members of the supervenience basis of B to members of the supervenience basis of A.

$P(A|B) = x$ where $0 \leq x \leq 1$.[16]

Conditional irrelevance is a screening off condition: due to the macrodependency B-A, conditionalising on B screens off the influence of LLD on A. In other words, in this case, all the influence of LLD (the microstate underlying B) is mediated through B.

These (in)dependencies can be represented graphically, as is familiar from both the literature on levels of dynamics and the causal modelling literature. Moreover, the conditional irrelevance condition is akin to the causal Markov condition in a multi-level setting: the most proximate causal parents of A screen off any other ancestor.

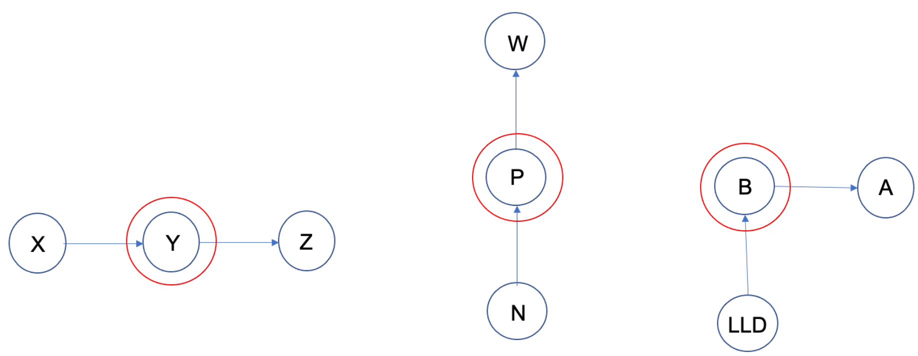

Figure 1: In all three diagrams, the red circled variable or fact screen off the downstream fact (Z, W or A respectively) from their ancestors: a generalised Markov condition.

The left hand diagram in Figure 1 shows a directed acyclic graph familiar from the causal modelling literature – the horizontal arrows represent causal dependence. The Causal Markov condition requires that nearest ancestors, i.e. more proximate causes, screen off more distant ancestors, such that Y screens off X from Z. In the grounding literature, see e.g. Schaffer (2016), there is an analogous Markov condition – thus, in the middle diagram the vertical arrows represent grounding. The P (psychological) facts screen off the N (neurophysical) facts from the W (well-being) facts. The interventionist framework generalises to a multilevel framework (Eva and Stern (2020)).

[16]Note this does not require that the macrodependency B-A is deterministic, i.e. $P(A|B) = 1$, if $x < 1$ then B could *still* screen off LLD.

In the right hand diagram, there is a multi-level Markov condition: B screens off the LLD from A.[17] The arrow from B to A represents the macrodependency, which could be causal, nomic or dynamical. We do not commit to the nature of the dependency from LLD to B, but it could just be taken to be supervenience. In this way 'unconditional relevance' fits with the supervenience physicalism of the non-reductive physicalist – but the arrows could represent other dependence relations such as grounding.

Now let's relax the assumption that the macrodependency B-A is strict. After all, it is far more realistic for special science laws to have exceptions, as Fodor (1974) emphasised. Special science laws, such as Mendel's law, the Law of demand and Snell's law, are frequently taken to be ceteris paribus laws for this reason.

Let's take a case of levels of dynamics, since this is the example that Fodor discusses. Exceptions to the macrodependency B-A occur when one (or more) members of the supervenience basis of B (the $LLD_i....LLD_N$ earlier) is not taken to the supervenience basis of A, as depicted in Figure 2. Such exceptions are sometimes called 'deviant microstates'.

In the case where the LLD is a deviant microstate, which member of the equivalence class is instantiated clearly *does* matter; so, in the example in Figure 2, LLD_3 would not raise the probability of A.

As shown in Figure 2, the microdynamics T take microstates from earlier to later. If we have a strict macrodependency, then B screens off LLD exactly. (So in terms of conditional irrelevance: $P(A|B\&LLD) = P(A|B)$.) But, as shown in the diagram, when not every member of the equivalence class of the LLD subvening B is taken to the supervenience basis of A, then we only have approximate screening off.

Consider the game of life (see e.g. Dennett (1991)): the microstate of the grid and the grid microdynamics are unconditionally relevant to the later macrostate, but are screened off by macrodynamics such as glider dynamics (a collection of cells that form a glider will travel with a certain velocity). But the glider dynamics in the game of life only hold until, e.g., a particularly per-

[17]For reasons that we will see shortly, this is an idealised case. When B-A is not a deterministic case of causation, then the multi-level Markov causation requires the subvening parent of A as well as A's causal parent, B for exact screening off, see Eva and Stern (2020) for more details.

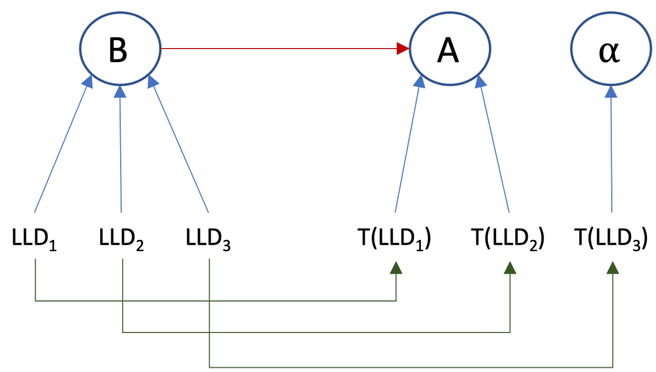

Figure 2: In this case, not every member of the supervenience basis of B is taken by the microdynamics T to the basis of A. This is how Fodor (1974) argues that the laws of special sciences have exceptions.

nicious piece of debris floats by destroying the glider pattern. Consequently, the glider dynamics only apply over certain timescales and for certain initial conditions.

Another example is found in statistical mechanics. There are two levels of dynamics: the reversible microdynamics, and the irreversible macrodynamics, such as the Boltzmann equation. The full microstates and microdynamics are unconditionally relevant to the later macrostate, but are screened off by the macrodynamics. There are, occasionally, deviant microstates that are not taken by T from the supervenience basis of B to the supervenience basis of A.

In both cases, the screening off condition is weakened to an approximate screening off. $P(A|B) \approx P(A|B\&LLD)$ because if the particular LLD is one of the deviant microstates such as LLD_3, then the $P(A|LLD_{dev})$ is very different from $P(A|LLD_{not-dev})$.

> **Conditional irrelevance (full generality):** $P(A|B\&LLD) \approx P(A|B) = x$ where $0 \leq x \leq 1$.

Why think that this still counts as 'screening off' if it's only approximate?

Even though it's approximate, it still has teeth: in particular, the approximate nature of the generalised screening off condition does not prevent it from delivering clear verdicts. The lower-level details are approximately screened off by statistical mechanical dependencies, which makes gases and entropy candidates for being emergent. Approximate screening off is, thus, essential for any account to deliver an ontology that matches that of the scientific realist; a strict criterion is not compatible with the nature of real-world relations.[18] But other putative dependencies (such as total cholesterol causes heart disease) do not even approximately screen off the lower-level details (whether it is high or low density cholesterol), and so total cholesterol is not a candidate emergent entity since the macrodependency in question does not screen off, even approximately.

If the approximate nature sits uneasily, there are two possible responses, both of which are independently used in the philosophy of physics and ceteris paribus laws literatures.

The first response is that background conditions may allow us to rule out the deviant microstates. In the statistical mechanical case, the system will approach equilibrium unless there are very finely balanced correlations in the microstate. These are ruled out in the choice of initial conditions used in constructing the macrodynamics from the microdynamics – this constraint on initial conditions is called 'Naturalness' by Wallace (2019). Arguably, this construction allows us to see the limited domain of applicability of the macrodynamics – they only apply for certain initial states. In these cases, by stipulating the right background conditions, the exact screening off condition is restored. That is, in examples such as statistical mechanics where we know the caveats, we can convert $P(A|B\&LLD) \approx P(A|B)$ to $P(A|B\&LLD) = P(A|B)$ by explicitly stating the deviant circumstances (e.g. the unnatural initial states, and timescales longer than the recurrence time).

More generally, this echoes certain approaches to understanding ceteris paribus laws – how to understand ceteris paribus laws turns out to be closely tied to one's metaphysics of laws (cf. Earman and Roberts (1999)).

For example, on Fodor (1991)'s account the ceteris paribus law is rendered strict by including extra information (known as 'completers') that rules out the deviant states or interfering factors (Pietroski and Rey (1995) take a similar view). For Lange (2002) and Woodward and Hitchcock (2003), unlike

[18]It's for this reason that we reject the assumption in J. Wilson (2010) that only strict elimination of degrees of freedom is ontologically relevant.

universal laws, ceteris paribus laws hold only for a limited range of counterfactual suppositions.

However, some approaches to the metaphysics of laws do not spell out the circumstances in which the ceteris paribus law/macrodependency holds, but just stipulate that under 'normal conditions' or 'with high probability' B-A (Schurz (2002)). This leads us to the second option for understanding the approximate screening off: namely, to note that in the case of macroscopic systems, the deviant microstates are very few and far between. Consequently, the approximate equality is incredibly close to being an equality sign. This response works especially well in statistical physics, where deviant microstates can be demonstrated to constitute a staggeringly small fraction of all possible microstates. Other special sciences will have their own standards for goodness of approximation.

To sum up: macrodependencies that screen off are key to our account. While we have focused on dynamical dependencies in explicating this idea, nomic and causal dependencies operate in the same way. An example of the former: the ideal gas law. The quantum state and dynamics are unconditionally relevant (to motivate this: if matter were not quantum, it wouldn't be stable), but the ideal gas law screens off these microdetails: the temperature is determined by pV/nR. Causal generalisations are used in biology. To briefly mention an example to which we'll return in section 4: the underlying amino acids are unconditionally relevant to the behaviour of a virus, but the macrodependency between the viral infection and the mottling of tobacco leaves screens off these lower-level details. The discussion of screening off is taken up again in section 5 where examples of screening off failure are discussed.

Note that our account is connected to autonomy in the following way: the macrodependency that screens off the LLD warrants the claim that the macrofact A is *autonomous*. For Fodor (1997), the autonomy of the special sciences is partially tied to the irreducibility of these descriptions. But our screening off condition, and the account of autonomy it comes from is compatible with inter-theoretic reduction. For instance, in the temperature example above, the higher-level variables and their inter-relations can be understood (or derived from) the lower-level descriptions in terms of molecular motion. Likewise in the case of autonomous macrodynamics, like the Boltzmann equation, the higher-level macrodynamics can be constructed from the underlying microdynamics. As such, theoretical reducibility is compatible with screening off. And yet all this conforms to the alternative characterisation of autonomy found in Fodor (1997, p. 160): "unimaginably complicated to-ings and fro-ings of bits and pieces at the extreme micro-level manage somehow to converge on stable macro-level properties".

3.2 Novelty

In many cases, the entities that feature in the macrodependencies that screen off are emergent. Viruses, phonons, gases, and wildebeest feature in macrodependencies that screen off and are emergent – they are archetypal inhabitants of the rainforest.

But consider this case: when two stars are in orbit around a common centre of mass, they may be considered to be a binary star system. Such a system will in general be orbiting further objects, such as, say, a super-massive black hole. The trajectories of the individual stars will be unconditionally relevant to the trajectory of the binary star system, but by conditionalising on their joint centre of mass, one can compute the trajectory of the binary star system around the black hole. So the dynamical dependency involving the centre of mass variable (e.g. Newton's equations) screens off the individual stars' trajectories. This is a case of screening off. But it seems wrong to say that the binary star system is admitted to the rainforest in addition to the two stars that compose it. Why?

The dependencies of binary star system are missing novelty relative to the underlying two stars, whereas gases and viruses are relatively novel in this respect. Thus, the binary star system isn't admitted to the rainforest because the nomic dependencies describing the binary star system, whether these are those of Newtonian gravitation or general relativity, are just the *same* nomic dependencies that describe the trajectories of the two individual stars. Distinct macrodependencies are required for novelty.

Our novelty criterion thus captures what Mitchell (2012) and Wimsatt (2007) are after with their stipulation that emergence is non-aggregative: some aggregates, especially those which involve simple averaging,[19] are not

[19]Further examples where simple averaging techniques are inapplicable in physics are discussed in Batterman (2013) and M. Wilson (2017).

described by distinct macrodependencies, but by the very same dependencies that describe their parts. While the invocation of non-simple aggregation can account for certain instances of distinct macrodependencies, this won't suffice for a fully general criterion since, e.g., the gas may be thought as an aggregate of its molecules yet gases are described by distinct macrodependencies from their molecules.

Distinct macrodependencies at different levels leads to cross-classification – that ontological divisions at a novel level of description cut across those at more fundamental levels; this underwrites the metaphysically substantive consequence of novelty. In sum: emergent entities have distinct modal profiles and can sustain distinct counterfactuals than their more or less fundamental counterparts – the discussion that follows captures this by cashing out the distinct-macrodependency criterion with respect to dynamical and causal dependencies respectively.

Among the criteria that real entities satisfy, novelty stands out as both the least controversial and the most contested: if an entity isn't novel in some sense then it's already included in the ontology at some other level, and including it would be doubling up, so everything in the ontology ought to have some novelty with respect to the remainder of the contents of the ontology. But it's difficult to find much consensus in the literature as to what novelty involves.[20]

Given our ontological focus, merely stating that novelty corresponds to epistemic criteria such as 'surprising relative to some reference class' or 'unexpected' (see e.g. Butterfield (2011a,b)) is insufficient.

One might hope, contra Butterfield, that more can be said than 'we know novelty when we see it'. Indeed, if there are no inter-subjective standards for novelty, then it's of very limited use as a criterion for emergence – the differing metaphysical tastes of scientists and analytic metaphysicians might impact their judgements about novelty, rendering the broader enterprise circular. Their differing tastes might themselves be explained by features of the judgments they have been exposed to. For instance, philosophers trained on logic judge that if the laws of fundamental physics are complete, then everything else 'follows from' this – not leaving much room for novelty. In

[20]See e.g. Butterfield (2011a), Chalmers (2006), Crowther (2015), Knox (2016), and J. Wilson (2010) for a range of distinct views.

contrast, physicists witness radically different behaviour in different domains. Fluid behaviour differs from atomic behaviour, quantum fields behave differently from ordinary particle mechanics. Most strikingly, the macroscopic behaviour described by classical mechanics differs radically from microscopic behaviour described by quantum mechanics. Thus, physicists would say that the classical world qualifies as novel, and so is emergent; see Wallace (2012).

For us, **a macrodependency is novel if and only if it's not type-identical to the microdependencies that instantiate it**. Note that this conception of novelty is inherently contrastive – a macrodependency is distinct with respect to a given microdependency.

The idea is that lower-level kinds or entities may be connected, joined up, or organised in various complex ways. As a consequence, the combined whole at the higher level participates in novel macrodependencies that are not identical with any microdependencies, and it's this which justifies the inclusion of that whole as a new kind in the rainforest.

One challenge faced by attempts to identify what's novel about higher-level entities is to avoid requiring too strong a species of novelty. While our account is compatible with higher level entities' being strongly emergent, strong emergence shouldn't be a foregone conclusion of our account. Instead, we aim to avoid the theoretical irreducibility criterion for higher-level ontology.

For us, novelty refers to distinct dependencies, but this distinctness manifests in slightly different guises for dynamical and causal dependencies – respectively, different functional form, and novel causal powers. In the remainder of this section, we develop this, and relate it to the concept of cross-classification.

What's crucial to us is that many of the theorists discussed in this section have made substantial progress in identifying what's required for ontological emergence, but that most have failed to identify precisely what's required. It's not sufficient that the higher-level description be surprising or that it be just a bit different from that at the lower level. Rather, we require that there be cross-classification and thus no identity between the types at the different levels. In the following, we demonstrate how many accounts get at this from slightly different perspectives, and that our general characterisation subsumes and improves on each.

3.2.1 Distinct Dynamical Dependencies

Taking dynamical dependencies first, an entity is novel if it features in macrodependencies with distinct functional form from the corresponding microdependencies. Gases, as described by thermodynamics, have a great many macrodependencies that are of a different type to the microdependencies that describe the behaviours of their particles. A well-studied example is that the particles may be described by time-reversible Hamiltonian dynamics, and the gas described by the time-irreversible Boltzmann equation. These clearly have a different functional form, and such functional forms give rise to different properties for each equation notwithstanding that, given certain conditions, one equation is derivable from the other.

This feature of distinct functional form is generic across physics: the relations between different theories involve differing dynamics and macrodependencies that aren't identical to any microdependencies. One further commonplace example is that the functional form of the dependencies of classical mechanics is distinct from those found in quantum physics, and this helps explain why both the classical and the quantum domains contribute entities to the rainforest.

3.2.2 Distinct Causal Dependencies

As should be clear from the preceding discussion, we think that the metaphysics of science should be understood in terms of dependencies. While dynamical and nomic dependencies are standard in physics, (see e.g., Price and Corry (2007) for arguments not to use causal language in physics), biological and higher-level sciences are often conceptualised in causal terms. We follow the relevant literature in cashing out novelty here in relation to causal powers, yet we are not committed to any particular conception of causal powers, and expect our approach to remain friendly to the Humean (cf. J. Wilson (2014)), In our view, where A depends on B, B may be understood to have the power or disposition to bring about A. The macrodependencies in which B participates determine what B can do, i.e. what powers B has. Given this connection, we can relate our discussion of novel macrodependencies to the literature on novel powers.

As further discussed in section 4, higher-level biological entities can do things that lower-level entities cannot do. For example, tobacco mosaic viruses (TMV) have the power to mottle tobacco leaves, whereas the amino acids that constitute the virus do not have that power. So, the dependencies between TMV and tobacco leaves are distinct and of distinct types from the dependencies between the amino acids found in the leaves and the virus. Such distinct macrodependencies have distinct functional forms, and thus, the same type of novelty is found here as in the physics case. Not only is this the same type of novelty as in the physics case, our conception of novelty may be understood to subsume the following approach.

J. Wilson (2011), in arguing that non-fundamental kinds can be causally efficacious, finds such efficacy not only with distinctive causal powers (which she sees as akin to strong emergence) but also "with distinctive collections of powers" (ibid. p.135).

The general idea, which has been taken up widely, see e.g. Tahko (2020), is that higher-level kinds feature a particular proper subset of the full range of causal powers located in the lower-level realising kinds. Novelty is then a consequence of the higher level having strictly fewer powers than its realisers. This guarantees that the higher level has a distinct but dependent causal profile. If an entity is individuated by its powers, then this characterisation underwrites additions to the ontology.

J. Wilson (2021, p. 51) draws an interesting distinction between 'fundamentally novel' and 'nonfundamentally novel' powers, where the former but not the latter entail strong emergence; only nonfundamentally novel powers are to be fully explained in terms of the subset strategy. In accordance with our overall aim to remain compatible with but not committed to theoretical reducibility, we remain neutral on whether our powers are fundamentally or nonfundamentally novel.[21] Given that proviso, a higher-level entity is novel if it has novel causal powers.

But we think that her account doesn't do enough to emphasise the importance of organisation/dynamical interaction for novel causal powers. Novelty is related to the dynamical structure, interactions, and organisation of parts in the entity – that's what's responsible for the new entity having strictly fewer powers than those of its parts, yet being able to do things that its parts

[21] Though we disagree with Wilson's connection between this distinction and aggregation.

cannot do individually.

3.2.3 Cross-classification

These notions are related to the more general concept of cross-classification. Wherever we find a genuinely novel macrodependency that's distinct from the corresponding microdependency then such dependencies will cut across one another. Consequently, the entities which we read off these dependencies are cross-classified.

For example, relative to the underlying proteins, viruses have novel causal powers. That is, they exhibit distinct macrodependencies associated with the following cross-classification: at the higher level the dependencies relate TMV to tobacco leaves, whereas at the lower-level dependencies relate, for example, the RNA to the proteins it codes for. What's significant for cross-classification is that some of the dependencies between the component parts of TMV are also found between the component parts of the tobacco leaf. So the ontological division into virus and leaf cuts across the ontological division into RNA and proteins. Likewise, in the physics case, the ontological division into gases and the containers (into which gases spread) cuts across the ontological division into molecules before and after collision. Sorting the entities into the higher-level kinds 'gas' and 'container' or 'leaf' and 'virus' won't line up with the lower-level's 'molecules with various speeds', 'proteins', and 'RNA'.

The macrodependencies responsible for screening off the microdetails are not type-identical to the microdependencies. And since we read off our ontology from these patterns/laws/dependencies, the cross-cutting of macro and microdependencies leads to a cross-classification of the ontological divisions or kinds at the higher and lower levels – cementing the notion that the higher-level entities are indeed distinct from the lower-level entities.

It's the novelty of the dependencies that lead to cross-classification and, together with the screening off of the dependencies, underwrites the status of the higher-level entities as ontologically emergent.

As noted above, this conception of novelty is compatible with theoretical reducibility. We can, for example, explain how proteins come together to form a virus that can do things the proteins can't do individually, or how collections of particles with certain initial conditions can act differently from the particles individually. Their interactions are governed by the microdependencies, but the totality is governed by the macrodependencies. That's not mysterious, and doesn't entail the denial of theoretical reductionism but, for the reasons discussed above, it does provide a basis for ontological emergence.

It's important to note that our focus on cross-classification is significantly inspired by Eleanor Knox's introduction of variable change as a criterion for emergence; see Franklin and Knox (2018) and Knox (2016). Changing variables is a useful strategy for uncovering dependencies that may turn out to be novel and screen off the microdetails.[22] In this way, Knox's account gives us the epistemic tools for finding novel macrodependencies that screen off – ontological emergence; but her account should not be understood as a purely epistemological account as some such as De Haro (2019) have read it.[23] Knox also emphasises a different kind of novelty to our account: she argues that the combination of changing variables and abstractions can give rise to novel explanatory value. For us this is because changing variables and abstractions uncover new dependencies, and these dependencies generate novel explanations by providing the causal or nomic link between explanantia and explananda. Our account thus underwrites the connection between novel explanatory value and emergence. Although we don't focus on explanation as a criterion for admittance into the rainforest, we embrace this consequence of our account.

Perhaps some will find this all to be inadequate: they might hold that real novelty requires a dependency relation that goes beyond what is *determined* by fundamental physics. So long as we maintain compatibility with theoretical reducibility to fundamental physics — or even just the supervenience between levels — then these folks will hold that no dependency counts as truly novel with respect to its reductive base.[24]

[22] Changing variables can help us find dependencies/patterns that are "there for the picking up if only we are lucky or clever enough to hit on the right perspective" Dennett (1991, p. 41).

[23] While we think this is a misreading, to avoid similar confusions about our project, we emphasise that variable changes are of especial importance when these help identify distinct dependencies and consequent cross-classification.

[24] Our response to such claims is to admit that this is indeed a salient philosophical understanding of 'novel'. However, we return to the claims expressed earlier: we hold that

To sum up, we think that theoretical reduction is widespread, but independent of the existence of non-fundamental entities. Rather, it's essential to those entities' existence that the world is differently arranged at different levels – that the best categories to describe nature at one level cut across those used at other levels. That's the basis of this view of novelty, and captures precisely the ontology implicit in the special science realist's project!

4 Case Study: A Virus

Pretty much everyone alive in 2020 or 2021 knows that viruses can have a causal impact both on individual humans, and, more specifically, on individual humans' lungs, olfactory systems, immune systems, and on human societies more generally. For those who think that to exist is to play an essential role in causally bringing about certain effects, or to exist is to play an essential role in scientific explanations, or to exist is to feature in scientific laws, it should be incontrovertible that viruses exist.[25] The case of viruses allows us to demonstrate that the criteria for inclusion in the rainforest set out above return an ontology that matches that of the scientific realist.

By establishing that the reasons for including viruses, or any other entity in our ontology do not relate to the failure of the theoretical reduction of such entities, we have made significant progress in one of the key aims of the paper – namely, establishing that theoretical reduction and ontological emergence are compatible. Moreover, the various examples above, especially the example of the gas as described by the Boltzmann equation also help establish this point.

Are viruses in fact theoretically reducible? Can one explain how viruses function by investigating their parts and the details of the interactions among their parts? To these questions we answer 'yes'.

Take, for example, the Tobacco Mosaic virus (TMV) – this is of course a

irreducibility is an inappropriate criterion for inclusion in the ontology. As such, given that novelty is the right kind of criterion, it's appropriate to interpret 'novelty' in a way that's compatible with reducibility. Thus, we hope to be judged by the upshots of our framework rather than falling prey to objections concerning the alternate meanings of individual terms.

[25]See Khalidi (2013, §5.4) for a discussion of *virus* as a natural kind.

very simple example of a virus – one could hold that only more complicated viruses exist over and above their parts, but this seems untenable, and TMV is well known to be virulent and have a devastating impact on tobacco crops – thus, it's difficult to challenge its causal profile.[26]

TMV's "rod shape results from its basic design, namely a regular helical array of identical protein subunits, in which framework is embedded a single molecule of RNA wound as a helix" (Klug (1999, p. 531)).

According to Liu and Nelson (2013): "Tobacco mosaic virus (TMV) encodes four known functional proteins: the 126 and 183 kDa replication-associated proteins, the movement protein (MP), and the structural capsid or coat protein (CP)." The coat protein, for example, can be further analysed into 2130 identical coat protein subunits following the right-handed helix of an accompanying RNA strand to produce a 300 nm hollow cylinder with an outer diameter of 18 nm and a 4 nm wide central channel.

Proteins are strings of amino acids. So there are many parts, but it's only when the parts conform to the appropriate structure that the virus is able to realise its function and that it thus has its novel explanatory and causal attributes.[27] The wrong assemblage of parts would be inert – viruses are able to infect hosts exactly when they are structured appropriately. When this happens abstractions from the detailed interactions of the parts (be that at the molecular or sub-molecular level) are legitimated and the dependencies that underwrite causal and explanatory claims are in terms either of the virus as a whole or in terms of the interactions of the proteins.

While there remain unanswered questions about the functioning of viruses and proteins, it seems both pessimistic and unwarranted to claim that such questions aren't answerable at all. At least we wouldn't bet against the vast array of research projects involved with mapping and understanding the complex details of how viruses work. As such, the theoretical reduction (according to some or other account) of viruses seems highly plausible.

We say: there are viruses in the rainforest (i.e. viruses feature in our

[26]While it's tempting to choose SARS-CoV-2, the science on TMV is, at the time of writing, far more settled.

[27]See Herman et al. (2021) for a discussion of structure and function in biology. Oppenheim and Putnam (1958) also discuss synthesising structurally identical but causally inert collections of proteins that mirror but are not identical to TMV.

best ontological inventory). Unless there is some radical and unforeseen scientific revolution, viruses are essential to our scientific account of the world. And, as such, those who accept that there are non-fundamental real kinds or entities ought to accept viruses into their ontology. That's the justification for including viruses in the ontology, but the explanation for viruses' inclusion over and above the parts which constitute viruses, is that viruses are ontologically emergent, that is: virus dependencies screen off lower-level descriptions and are novel.

Virus dependencies screen off: the molecular and submolecular descriptions are unconditionally relevant, but in part conditionally irrelevant to the virus's efficacy. The dependency between a viral infection and, say, the mottling of tobacco leaves screens off the sub-parts of the virus.

The presence of an active TMV will be sufficient to enable predictions regarding whether or not tobacco leaves will mottle or otherwise register signs of infection. As such, conditional on a property of the virus as a whole – whether it is active or inactive – one can abstract away from the lower-level details that compose the virus: the macrodependencies about the virus screen off these lower-level details.

Virus dependencies are novel: as discussed above they have novel causal powers – they can do things their parts can't do; and virus-leaf dependencies are distinct from the protein and amino acid dependencies – while the former engender ontological divisions into plants and viruses, the latter engender much more fine grained divisions into types of protein or amino acid, which can be found both in viruses and plants.

Now it could be the case that TMV is like jade[28] and that in fact there are TMV sub-types that lead to different kinds of mottling and have different structures. If that did turn out to the best analysis of TMV its dependencies would fail our novelty criterion as TMV would have no novel causal powers relative to its sub-types. Consider by analogy the claim that SARS-CoV-2 causes Covid vs. the suggestion that variants of the virus cause variants of the disease. This dependence of our ontology on scientific discovery is exactly as it should be.

[28]See Franklin (2021) and Kim (1992).

5 Taking Stock: Trout-Turkeys and Other Philosophical Conundra

The previous section showed us that viruses are admitted to the rainforest, now it's worth demonstrating that some putative entities are denied entry. We are after a verdant and lush rainforest, but gerrymandered kinds ought not to be there!

To be clear: we do not claim that the contents of the rainforest exhaust all our ontology. As such, we take no stance as to the existence of mathematical entities, or gerrymandered kinds such as trout-turkeys, or any other ontology. Rather we are interested in specifying admission criteria to the rainforest, and we do claim that the rainforest has a particular role to play since it contains all and only those entities to which the scientific realists ought to be committed.[29]

So, are there trout-turkeys in the rainforest?[30] The answer is, obviously, 'no', but to see precisely why we ought to specify a bit further exactly what trout-turkeys are.

Trout-turkeys could be conceived of as mereological sums of trout and turkeys. That is, each trout-turkey is composed of one specific trout and one specific turkey, whether or not these animals happen to be spatially contiguous at any particular time.[31]

Are trout-turkeys emergent? First we need to consider the screening off condition. Straightaway the situation is different: in the other cases, there were bona fide scientific theories at both the lower and higher level. What are the two levels for a trout-turkey? The higher-level theory putatively describes the mereological complex 'a trout-turkey' and the lower-level theory must either describe the collection of simples (say, atoms) that compose the trout-turkey, or perhaps the trout and the turkey.

[29]If truly fundamental entities were ever to be discovered then these would form a notable exception to this claim as these would not be ontologically emergent but would feature in the scientific realist's ontology.

[30]There is a separate question (the new problem of induction) about what, if anything, makes trout (or green) a better variable than trout-turkey (or grue), but we don't have space to address that here.

[31]If, instead trout-turkeys were composed of half a particular trout and half a particular turkey, the claims in the following paragraph would still apply.

To state the obvious, no scientific theory of trout-turkeys is forthcoming. Do they swim? Are they prey of other animals? How do they reproduce? But is this lack of scientific theory for them just because of the basic kinds and predicates *we* choose to describe the world? Could there be some kind of creature that does choose to use the variable 'trout-turkey' and has higher-level theories about their evolution? While of course the question of variable choice is a tricky one entangled with the problem of induction (see e.g. Cohen and Callender (2009)), we have principled reasons to think that trout-turkeys are not emergent.

To demonstrate this, it is helpful to remember other examples of gerrymandered kinds, like Dennett's lost socks, or the centre of mass of all philosophers of science. Let's assume we are interested in how the location of these kinds changes over time. What description can we give for these entities? Well, we can see how their components' locations change over time, and use that to see how the composite object changes location over time. Are there dependencies involving these kinds that screen off lower-level details? No! In order to see how the centre of mass changes at each time step we have to look at where each member has moved to and re-average – there is no higher-level equation/macrodependency that screens off these microdetails.[32] The locations of the components remain relevant.

Following M. Wilson (1985), one could come up with a mathematical function to describe any arbitrary concatenation of entities, but this wouldn't take the form such that it would allow for projecting into the future. So, it's not at all clear that this should count as a macrodependency at all; this point is made by Ross (2000, p. 162), albeit in a different idiom. More importantly, this macrodependency would not screen off any lower-level details for the reasons given above: to know what our trout-turkey will do, we have to check what the trout and what the turkey do, and then re-sum. In other words: there is no conditional irrelevance, and so the screening off condition is not met. Trout-turkeys are not emergent.

One might have the worry at this stage that our emergent ontology is irrecoverably vague and that this is a reason to reject the approach outright. In particular, one might worry that conditional irrelevance was not an exact, cut and dried condition, but involved *approximate* screening off. Oftentimes,

[32]Perhaps there are certain higher-level regularities about the centre of mass of all philosophers of science – the location shifts whenever there is a PSA or BSPS!

this approximation can be made exact by specifying the exact circumstances when the lower-level details become relevant again. But when we can't do this, should we be concerned about the vagueness that creeps in?

We think not. Consider clouds: not only do these have blurry edges, but there will be a vagueness as to whether a collection of water droplets counts as a cloud or not. We don't think that such vagueness can be excised, in part because we can't think of any scientific examples that don't have vague boundaries and vague edge cases. That is, even in fundamental physics, particles are excitations of quantum fields, and there's a vagueness both as to the boundaries of such excitations, and when an excitation counts as a particle or not. Our goal in this paper is not to persuade the desert advocate – if one has the view that future, as-yet-undiscovered fundamental physics will furnish a complete and precise ontology then that's fine. But as long as one is prepared to admit some non-fundamental entities, some degree of vagueness in ontology is inevitable.

Having addressed these concerns we come to the consideration of how our ontology relates to the well-known articulations of rainforest realism in the existing literature.

6 Rainforest Realism

Our concern in this paper has been establishing what lives in the many-levelled ontology of the scientific realist.[33] The terminology of 'rainforest realism' comes from Ladyman and Ross (2007) and Ross (2000); Ladyman (2017, p. 152) identifies this with the claim that "[t]he ontological commitments of the special sciences should be taken as metaphysically on a par with those of physics". So how does our account of the higher-level, or special science ontology relate to Ladyman and Ross's? First, we discuss how our account aligns with theirs, and then we discuss two ways in which our account is less radical.

Dependencies play a key role in our analysis, and are clearly patterns of

[33]Note that Ladyman and Ross (2007) and Shech and McGivern (2019) call this ontology 'scale relative'. While, broadly speaking, this fits with our conception of the rainforest, we don't use this terminology because levels, as understood in this paper, don't always align with particular scales, cf. Potochnik (2017).

some kind. But the link is tighter than this: the novel macrodependencies that screen off lower-level details are *real patterns*. In this way, our account offers an alternative explication of how patterns at one level relate to patterns at another level such that there is a lush ontology. And our account needn't be considered a rival to Dennettian real patterns – indeed Ladyman explicitly allows that the "theory of Real Patterns can be explicated in various ways" (Ladyman (2017, claim XIV)).

Whereas the account in Ladyman and Ross (2007) focusses more on the features of scientific representations, we emphasise the worldly backing of such features. In addition both we and they (ibid. p.233) insist on the projectibility of real patterns and dependencies: that is, that they must play a role in scientific generalisations that can serve as the basis for inductive reasoning. Likewise, where they discuss the compressibility of an encoding relative to the bitmap (see Dennett (1991)), our notion of screening off underlies such compressibility. Thus, our account of the rainforest encapsulates significant features of Ladyman and Ross's project.

One way in which our account is less radical than theirs is the status we give to dependencies and patterns. For Ladyman and Ross, patterns are ontologically prior to the entities that enter into them. Indeed, more radical views, such as French (2014), go further and argue there is only relational structure and all entities should be eliminated. While our account places dependencies and patterns centre stage, and is compatible with structural realism, it does not require or presuppose structural realism. On our account, entities ontologically emerge only if they are involved in – featured on one side or the other – of some dependency that is both novel and screens off lower-level details. It's the relations between dependencies discussed in detail above, in virtue of which entities emerge. But the primacy of patterns or dependencies in our account is not a metaphysical claim: it is just that patterns are the starting place for our analysis and so are conceptually prior – this needn't entail that they are ontologically prior.

There is a further way in which our account can be considered to be a generalisation of Ladyman and Ross: our language of autonomy and novelty generalises away from, and subsumes, the information-theoretic gloss that Ladyman and Ross give. Information-theoretic ideas are useful ways to express the key ideas and connect the philosophical discussion to empirical research in data science, but they are eliminable.

Ladyman and Ross defend their view by claiming that it “expresses the principle of Occam’s razor, in restricting ontological commitment to what is *required* for a maximally empirically adequate science”, and they spell this out in terms of predictive novelty (ibid., pp. 233-234, original emphasis). We differ from them by focussing on a slightly different form of novelty: even though higher-level goings-on may be predictable from the bottom up, the novel dependencies are not type-identical with those at the lower level. By tying the definition of real patterns to the availability of novel and screening-off dependencies at higher levels, we satisfy the intuition that such dependencies are required by maximally empirically adequate science.

On our view, real patterns might be both reducible and emergent. While some reductions may engender elimination, those patterns which are emergent are ineliminable. Even if it turns out that reduction is commonplace, this will allow for a verdant rainforest, with an array of entities whose properties may be understood from the bottom up.

7 Conclusion

Ontological tastes differ: some prefer lush rainforests and others stark deserts. While we do not presume to have persuaded those with the latter tastes, we hope to have articulated an account that makes explicit the commitments popular among philosophers of science. Our goal is to admit some non-fundamental scientific entities into the ontology, and to exclude gerrymandered non-scientific kinds. We claimed that there are principled ways to make that distinction.

We argued that those entities that are to be admitted to the rainforest are ontologically emergent, and yet at least some of them may be theoretically reducible. We went on to characterise ontologically emergent entities as those that participate in novel and screened-off dependencies. Through a careful application of these criteria, we demonstrated that viruses are in, but that trout-turkeys are out. Finally, we showed that this approach sets rainforest realism on firmer metaphysical ground.

References

Barnes, Elizabeth (2012). "Emergence and Fundamentality". In: *Mind* 121.484, pp. 873–901. ISSN: 00264423, 14602113. URL: http://www.jstor.org/stable/23407311.

Batterman, Robert W. (2013). "The Tyranny of Scales". In: *The Oxford Handbook of Philosophy of Physics*. Ed. by Robert W. Batterman. Oxford University Press, pp. 256–286.

Bedau, Mark A and Paul Ed Humphreys (2008). *Emergence: Contemporary readings in philosophy and science.* Cambridge, Mass: MIT press.

Bennett, Karen (2007). "Mental causation". In: *Philosophy compass* 2.2, pp. 316–337.

Bogen, James and James Woodward (1988). "Saving the phenomena". In: *The Philosophical Review* 97.3, pp. 303–352.

Butterfield, Jeremy (June 2011a). "Emergence, Reduction and Supervenience: A Varied Landscape". In: *Foundations of Physics* 41.6, pp. 920–959. DOI: 10.1007/s10701-011-9549-0.

— (June 2011b). "Less is Different: Emergence and Reduction Reconciled". In: *Foundations of Physics* 41.6, pp. 1065–1135. DOI: 10.1007/s10701-010-9516-1.

— (2014). "Reduction, Emergence, and Renormalization". In: *The Journal of Philosophy* 111.1, pp. 5–49.

Chakravartty, Anjan (2017). "Scientific Realism". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2017. Metaphysics Research Lab, Stanford University.

Chalmers, David J. (2006). "Strong and Weak Emergence". In: *The Re-Emergence of Emergence: The Emergentist Hypothesis From Science to Religion*. Ed. by P. Davies and P. Clayton. Oxford University Press.

Cohen, Jonathan and Craig Callender (2009). "A better best system account of lawhood". In: *Philosophical Studies* 145.1, pp. 1–34.

Crane, Tim (2001). "The Significance of Emergence". In: *Physicalism and its Discontents*. Ed. by Barry Loewer and Grant Gillett. Cambridge, UK: Cambridge University Press.

Crowther, Karen (2015). "Decoupling emergence and reduction in physics". In: *European Journal for Philosophy of Science* 5.3, pp. 419–445.

De Haro, Sebastian (2019). "Towards a theory of emergence for the physical sciences". In: *European Journal for Philosophy of Science* 9.3, pp. 1–52.

Dennett, Daniel C. (1991). "Real patterns". In: *The Journal of Philosophy*, pp. 27–51.

Dizadji-Bahmani, Foad, Roman Frigg, and Stephan Hartmann (2010). "Who's Afraid of Nagelian Reduction?" In: *Erkenntnis* 73.3, pp. 393–412.

Dorr, Cian and John Hawthorne (2013). "1. Naturalness". In: *Oxford Studies in Metaphysics, Volume 8*, p. 1.

Earman, John and John T. Roberts (1999). ""Ceteris Paribus", There Is No Problem of Provisos". In: *Synthese* 118.3, pp. 439–478. DOI: 10.1023/a:1005106917477.

Eva, Ben and Rueben Stern (2020). "Antireductionist Interventionism". In: *British Journal for the Philosophy of Science*.

Fairchild, Maegan (Apr. 2019). "The Barest Flutter of the Smallest Leaf: Understanding Material Plenitude". In: *The Philosophical Review* 128.2, pp. 143–178. ISSN: 0031-8108. DOI: 10.1215/00318108-7374932. eprint: https://read.dukeupress.edu/the-philosophical-review/article-pdf/128/2/143/575793/143fairchild.pdf. URL: https://doi.org/10.1215/00318108-7374932.

Fodor, Jerry A. (1974). "Special Sciences (Or: The Disunity of Science as a Working Hypothesis)". In: *Synthese* 28.2, pp. 97–115.

— (Jan. 1991). "You Can Fool Some of The People All of The Time, Everything Else Being Equal; Hedged Laws and Psychological Explanations1". In: *Mind* C.397, pp. 19–34. ISSN: 0026-4423. DOI: 10.1093/mind/C.397.19. eprint: https://academic.oup.com/mind/article-pdf/C/397/19/10888930/19.pdf. URL: https://doi.org/10.1093/mind/C.397.19.

— (1997). "Special sciences: Still autonomous after all these years". In: *Noûs* 31.s11, pp. 149–163.

Franklin, Alexander (2021). "Can Multiple Realisation Be Explained?" In: *Philosophy* 96 (1). DOI: 10.1017/S0031819120000285.

Franklin, Alexander and Eleanor Knox (2018). "Emergence without limits: The case of phonons". In: *Studies In History and Philosophy of Modern Physics* 64, pp. 68–78. DOI: 10.1016/j.shpsb.2018.06.001.

French, Steven (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press.

Gillett, Carl (2016). *Reduction and emergence in science and philosophy*. Cambridge University Press.

Hellman, Geoffrey Paul and Frank Wilson Thompson (1976). "Physicalism: Ontology, determination, and reduction". In: *The Journal of Philosophy* 72.17, pp. 551–564.

Hendry, Robin Findlay (2010). "Emergence vs. reduction in chemistry". In: *Emergence in mind*, pp. 205–221.

Herman, MA et al. (2021). "A Unifying Framework for Understanding Biological Structures and Functions Across Levels of Biological Organization". In: *Integrative and Comparative Biology* 61.6, pp. 2038–2047.

Hoefer, Carl and Genoveva Martí (2020). "Realism, reference & perspective". In: *European Journal for Philosophy of Science* 10.3, pp. 1–22.

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

Kim, Jaegwon (1992). "Multiple realization and the metaphysics of reduction". In: *Philosophy and phenomenological research* 52.1, pp. 1–26.

Klug, Aaron (1999). "The tobacco mosaic virus particle: structure and assembly". In: *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 354.1383, pp. 531–535.

Knox, Eleanor (2016). "Abstraction and its Limits: Finding Space For Novel Explanation". In: *Noûs* 50.1, pp. 41–60. DOI: 10.1111/nous.12120.

Ladyman, James (2017). "An Apology for Naturalized Metaphysics". In: *Metaphysics and the Philosophy of Science: New Essays*. Ed. by Matthew Slater and Zanja Yudell. Oxford University Press. Chap. 7. DOI: 10.1093/acprof:oso/9780199363209.003.0008.

Ladyman, James and Don Ross (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press. With David Spurrett and John Collier.

Lange, Marc (2002). "Who's afraid of ceteris-paribus laws? Or: How I learned to stop worrying and love them". In: *Erkenntnis* 57.3, pp. 407–423.

Liu, Chengke and Richard S Nelson (2013). "The cell biology of Tobacco mosaic virus replication and movement". In: *Frontiers in plant science* 4, p. 12.

Mitchell, Sandra D. (2012). "Emergence: Logical, functional and dynamical". In: *Synthese* 185.2, pp. 171–186.

Nagel, Ernest (1961). *The Structure of Science: Problems in the Logic of Scientific Explanation*. Harcourt, Brace & World, Inc.

Ney, Alyssa (2021). "Reductionism". In: *The Internet Encyclopedia of Philosophy*. URL: https://iep.utm.edu/red-ism/.

Oppenheim, Paul and Hilary Putnam (1958). "Unity of Science as a Working Hypothesis". In: *Minnesota Studies in the Philosophy of Science* 2, pp. 3–36.

Pietroski, Paul and Georges Rey (1995). "When Other Things Aren't Equal: Saving Ceteris Paribus Laws from Vacuity". In: *The British Journal for the Philosophy of Science* 46.1, pp. 81–110. ISSN: 00070882, 14643537. URL: http://www.jstor.org/stable/687634.

Potochnik, Angela (2017). *Idealization and the aims of science*. Chicago: University of Chicago Press.

Price, Huw and Richard Corry, eds. (2007). *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*. Oxford University Press.

Robertson, Katie (n.d.). "Autonomy generalised; or, Why doesn't physics matter more?" In: *Ergo* (). Forthcoming.

Rosaler, Joshua (2017). "Reduction as an *A Posteriori* Relation". In: *The British Journal for the Philosophy of Science*. DOI: 10.1093/bjps/axx026.

Ross, Don (2000). "Rainforest realism: A Dennettian theory of existence". In: *Dennett's Philosophy: A Comprehensive Assessment*. Ed. by Don Ross, Andrew Brook, and David Thompson. MIT Press.

Schaffer, Jonathan (2004). "Two conceptions of sparse properties". In: *Pacific Philosophical Quarterly* 85.1, pp. 92–102.

— (2016). "Grounding in the image of causation". In: *Philosophical Studies* 173.1, pp. 49–100.

Schaffner, Kenneth F. (2013). "Ernest Nagel and Reduction". In: *The Journal of Philosophy* 109.8/9, pp. 534–565.

Schurz, Gerhard (2002). "Ceteris Paribus Laws: Classification and Deconstruction". In: *Erkenntnis (1975-)* 57.3, pp. 351–372. ISSN: 01650106, 15728420. URL: http://www.jstor.org/stable/20013162.

Seifert, Vanessa A. (2020). "The strong emergence of molecular structure". In: *European Journal for Philosophy of Science* 10.3, p. 45.

Shech, Elay and Patrick McGivern (2019). "Fundamentality, Scale, and the Fractional Quantum Hall Effect". In: *Erkenntnis*, pp. 1–20.

Sober, Elliott (1999). "Physicalism from a probabilistic point of view". In: *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 95.1/2, pp. 135–174.

Spirtes, Peter, Clark N Glymour, and Richard Scheines (2000). *Causation, prediction, and search*. MIT press.

Tahko, Tuomas E (2020). "Where do you get your protein? Or: biochemical realization". In: *The British journal for the philosophy of science* 71.3, pp. 799–825.

Taylor, Elanor (2015). "An explication of emergence". In: *Philosophical Studies* 172.3, pp. 653–669.

Thomasson, Amie L. (2009). "The Easy Approach to Ontology". In: *Axiomathes* 19.1, pp. 1–15. DOI: 10.1007/s10516-008-9057-9.

Wallace, David (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford University Press.

— (2019). "Naturalness and emergence". In: *The Monist* 102.4, pp. 499–524.

Waters, C Kenneth (2008). "Beyond Theoretical Reduction and Layer-Cake Antireduction: How DNA Retooled Genetics and Transformed Biological Practice". In: *The Oxford Handbook of Philosophy of Biology*. Ed. by Michael Ruse. Oxford University Press.

Wilson, Jessica (2010). "Non-reductive Physicalism and Degrees of Freedom". In: *The British Journal for the Philosophy of Science* 61.2, pp. 279–311. DOI: 10.1093/bjps/axp040.

— (2011). "Non-reductive realization and the powers-based subset strategy". In: *The Monist* 94.1, pp. 121–154.

— (2014). *Metaphysical Emergence: Weak and Strong*. URL: http://philsci-archive.pitt.edu/11200/.

— (2021). *Metaphysical emergence*. Oxford University Press, USA.

Wilson, Mark (1985). "What is This Thing Called 'Pain'?" In: *Pacific Philosophical Quarterly* 66.1.

— (2017). *Physics Avoidance: Essays in Conceptual Strategy*. Oxford University Press.

Wimsatt, William C. (2007). *Re-engineering philosophy for limited beings: Piecewise approximations to reality*. Harvard University Press.

Woodward, James (2021). "Explanatory autonomy: the role of proportionality, stability, and conditional irrelevance". In: *Synthese* 198.1, pp. 237–265. DOI: 10.1007/s11229-018-01998-6.

Woodward, James and Christopher Hitchcock (2003). "Explanatory generalizations, part I: A counterfactual account". In: *Noûs* 37.1, pp. 1–24.