

Emerging into the Rainforest: Emergence and Special Science Ontology

Alexander Franklin* and Katie Robertson‡

November 2021

Abstract

Many philosophers of science are ontologically committed to a lush rainforest of special science entities (Ross (2000)), but are often reticent about the criteria that determine which entities count as real. On the other hand, the metaphysics literature is much more forthcoming about such criteria, but often links ontological commitment to irreducibility. We argue that the irreducibility criteria are in tension with scientific realism: for example, they would exclude viruses, which are plausibly theoretically reducible and yet play a sufficiently important role in scientific accounts of the world that they should be included in our ontology.

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 trouts and turkeys, but philosophically gerrymandered objects like trout-turkeys do not qualify.

*alexander.r.franklin@kcl.ac.uk

†k.e.robertson@bham.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. 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 final theory of physics.

A common analogy relates 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 this is in tension with scientific realism: when debating whether viruses exist, the concern was 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 of quantum gravity.¹

But the desert is not to everyone's tastes. At the other extreme, some (such as Schaffer (2010)) dine on an 'ontological free lunch': all the non-fundamental entities such as mammals, viruses and mountains are part of their ontology for no extra cost.² This free lunch is a banquet: not only does it include viruses and glaciers but also gerrymandered kinds such as trout-turkeys and concatenations of spacetime regions. 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. But these gerrymandered kinds do not sit easily with the scientific realist. If ontological commitment is so easily had, 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

¹See Hoefer and Campillo (forthcoming) for an argument that the scientific realist should be committed to the entities of all sciences *other* than fundamental physics.

²In this paper we are neutral on the second-order metametaphysical question as to whether dependence or ontology is the key concern of metaphysics; we are interested in the first-order question of which things are real and which are not.

Goldilocks, she needs it to be ‘just right’: trouts, 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 in our accounts that are discussed later, there are also significant divergences, the foremost being the role that reduction plays.³

Traditionally, the existence of special science entities has hinged upon whether they are ‘reducible’ or not; if it should turn out that a special science kind is reducible then it ought to be eliminated from the rainforest.

We argue against this view. We say that viruses exist notwithstanding that we can understand them in terms of lower-level entities, such as nucleotides. Even if viruses, phonons, and mountains can be understood entirely from the bottom up (that is, they are reducible), they still roam the scientific realist’s rainforest – provided that they fulfil a certain condition: they must be emergent, by our lights. Provided that a non-fundamental, or special science entity is emergent, it is inoculated against being eliminated by any theoretical reduction.

In this paper, we will give a new account of emergence that is compatible with reduction. Then we will claim that the higher-level ontology is *emergent* in this sense: viruses, gases, and phonons are emergent unlike trout-turkeys and concatenations of random spacetime regions. In this sense, the entities of the special sciences emerge into a lush rainforest. Thus, we set out a novel metaphysical view that should be compelling for the scientific realist, and as such, central to the project of 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. Nonetheless, reduction is important to categorising types of emergence. Section 3 outlines our account of emergence. According to our account, 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

³In this paper we do not aim to argue against the view of Schaffer, or material plenitude in general. All of the gerrymandered kinds may exist, or be real in some sense, but they are importantly different from the kinds that the scientific realist is committed to. Our aim is just to spell out the difference between the entities that live in the scientific realist’s rainforest and the gerrymandered kinds.

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, but avoids some of the criticism levelled against their project. In section 7, we conclude.

2 The 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 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* just 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 interest to those who are realist 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_i

⁴See also e.g. Gillett (2016).

⁵The 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.

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 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. This traditional view is, in fact, endorsed by Nagel (1961, p. 337) who says that “successful [inter-theoretic] reduction of thermodynamics to statistical mechanics is [...] 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

⁶There are other related accounts around; see e.g. Dizadji-Bahmani, Frigg, and Hartmann (2010), Rosaler (2017), and Schaffner (2013).

⁷Nagel'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.

⁸Even those who advocate strong emergence, such as Hendry (2010), aim to maintain

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* 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. But 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. Later we will see that the case study of viruses is 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.

In this paper we develop an account of ontological emergence that 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.

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

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.

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

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.

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, as will be discussed in section 3.1. Secondly, the macrodependencies must be suitably novel, as will be discussed in section 3.2.

But what are the macrodependencies that we have in mind? Macrodependencies are projectible; each macrodependency may fall into various categories, 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 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)), and clearly accounts that define emergence as involving ‘surprise’ or ‘unexpected features’ render emergence a feature of our knowledge of the world. But we think that

emergence is 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 theories – our scientific representations of the world. 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 with the standard scientific realist, we presume that there are higher-level laws, higher-level causal relations, and higher-level dynamics.¹⁰

In general, emergence balances some form of dependence on the more fundamental with some form of independence, autonomy or distinctness.¹¹ Our account of emergence is slightly unorthodox in that we won’t have one 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’.

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

¹⁰The 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)).

¹¹Taylor (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”.

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 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 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 circumstances¹², $P(A|LLD) = 1$.

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 influence macrofact A .

¹²If the microdynamics take all the members of the supervenience basis of B to members of the supervenience basis of A.

Now onto 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 $100^{\circ}C$), 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:

$$\text{Conditional irrelevance (strict version): } P(A|B\&LLD) = P(A|B) = x \text{ where } 0 \leq x \leq 1.^{13}$$

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 screens off any other ancestor.

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)).

¹³Note 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.




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

In the right hand diagram, there is a multi-level Markov condition: B screens off the LLD from A.¹⁴ 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 $LLD_i \dots LLD_N$ earlier) is not taken to the supervenience basis of A, as depicted in Figure 2. Such exceptions are sometimes called ‘deviant microstates’.

¹⁴For 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 sub-vening parent of A as well as A’s causal parent, B for exact screening off, see Eva and Stern (2020) for more details.

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, LLD₃ would not raise the probability of A.

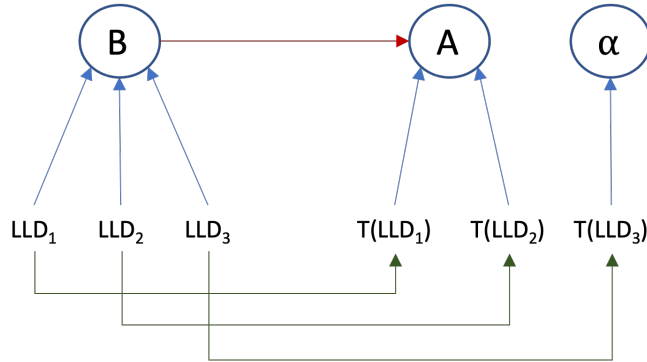


Figure 2: In this case, not every member of the supervenience basis of B is taken by the microdynamics T to the basis of A. This is how Fodor (1974) argues that the laws of special sciences have exceptions.

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 pernicious 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 $P(A|LLD_{notdev})$.

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. 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 literature, and *ceteris paribus* laws literature where examples are often taken from the special sciences.

The first response is that background conditions may allow us to rule out the deviant microstates as follows. 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’s stability account of laws and Woodward and Hitchcock’s invariance account, unlike 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 that will be discussed in detail 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.

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.

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 us to say that the binary star system is admitted to the rainforest in addition to the two stars that compose it.

We claim that the dependencies of binary star systems 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. Our novelty criterion thus captures what Mitchell (2012) and Wimsatt (2007) are after with their stipulation that emergence is non-aggregative: an aggregate of arbitrary parts doesn't *emerge* unless some novel macrodependency is identified.

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.¹⁵

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 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 novel with respect to a given microdependency.

¹⁵See e.g. Butterfield (2011a), Chalmers (2006), Knox (2016), and J. Wilson (2010) for a range of distinct views.

The idea is that lower-level kinds or entities may be connected, joined up, and organised in a multitude of ways. As a consequence of these organisations 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. Those who regard 'novelty' as essentially incompatible with theoretical reduction may substitute 'distinctness' throughout.

Novelty manifests in slightly different guises in physics and biology – 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.

Taking physics 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.

As should be clear from the preceding discussion, we think that the metaphysics of science should be understood in terms of dependencies. While dy-

namical 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-reductive 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. Given that proviso, **a higher-level entity is novel if it has novel causal powers.**

But we think that there are two extensions to her account that our approach can give. Firstly, we emphasise the importance of organisation/dynamical interaction for novel causal powers, and secondly and relatedly, we connect this to cross classification.

On the first point, it's important that 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 cannot do individually.

On the second point, these notions are connected by the more general concept of cross-classification. This is the idea that ontological divisions at one level are distinct from the ontological divisions at another level. Wherever we find a genuinely novel macrodependency that's sufficiently 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 on-

tology 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, thus, 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 microphysicalism, 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 Knox (2016). Changing variables is a useful strategy for uncovering dependencies that may turn out to be novel and screen off the microdetails.¹⁶ In this way, our account elaborates Knox's account as a metaphysical account of emergence, rather than as an epistemological account as some, such as De Haro (2019) have read it. 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. Knox also emphasises a different kind of novelty to our account: she argues that the combination of change of variables and abstractions gives 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.

¹⁶Changing 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).

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, then these folks will hold that no dependency counts as truly novel with respect to its reductive base.

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 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.

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.¹⁷ If ‘ontologically emergent’ just means exists and is non-fundamental, then viruses are, therefore, ontologically emergent.

As such, in this section the case of viruses is presented as a useful exemplar of a kind (or many kinds) that certainly should be taken to exist, but may well be understood from the bottom up, and, at least, insofar as they can’t be understood entirely from the bottom up, that’s not the reason that they exist.

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’.

¹⁷See Khalidi (2013, §5.4) for a discussion of *virus* as a natural kind.

Take, for example, the Tobacco Mosaic virus (TMV) – this is of course a 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.¹⁸

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 are arranged in the right way into the whole that the virus gets its novel explanatory and causal attributes.¹⁹ The wrong assemblage of parts would be inert – viruses are able to infect hosts exactly when their parts are assembled in the right way. When this happens abstractions from the detailed interactions of the parts 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 best ontological inventory). Unless there is some radical and unforeseen scientific

¹⁸While it’s tempting to choose Sars-Cov-2, the science on TMV is, at the time of writing, far more settled.

¹⁹Oppenheim and Putnam discuss synthesising structurally identical but causally inert collections of proteins that mirror but are not identical to TMV.

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: that virus dependencies screen off lower-level descriptions and are novel.

Virus dependencies screen off: the behaviour of the amino acids is unconditionally relevant, but 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 any details of the state of the individual proteins that compose the virus, because 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.

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.²⁰

So, are there trout-turkeys in the rainforest? 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 trouts 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.²¹

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.

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 to make up an example, the centre of mass of all philosophers of science. Let’s assume we are interested

²⁰If 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.

²¹If, instead trout-turkeys were composed of half a particular trout and half a particular turkey, the claims in the following paragraph would still apply.

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' location changes 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.²² 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, 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, does an irking vagueness creep 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

²²Perhaps 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 BPS!

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.²³ The terminology of 'rainforest realism' comes from Ladyman and Ross (2007) and Ross (2000), so how does our account of the higher-level, or special science ontology relate to theirs? 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 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. While we both (ibid. p.233) insist on the projectibility of real patterns and dependencies, where they discuss the compressibility of an encoding relative to the bitmap, our notion of screening off underlies such compressibility. Thus, our account of the rainforest encapsulates significant features of Ladyman and Ross's project.

²³Note 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).

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.

Acknowledgements

We are very grateful to participants in the FraMEPhys and MetaScience reading groups for their engagement and helpful feedback. We'd also like to thank audiences at MetaScience's 'Going Up' conference, and the EPSA 2021 conference. KR's work was supported by the FraMEPhys ERC grant 757295 and a Leverhulme Early Career Fellowship.

References

- Barnes, Elizabeth (2012). "Emergence and Fundamentality". In: *Mind* 121.484, pp. 873–901. ISSN: 00264423, 14602113. URL: <http://www.jstor.org/stable/23407311>.
- 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.
- 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.
- 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.
- 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>.
- Fodor, Jerry A. (1997). “Special sciences: Still autonomous after all these years”. In: *Noûs* 31.1, pp. 149–163.
- 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.
- Hofer, Carl and Genoveva Martí Campillo (forthcoming). “Realism, Reference and Perspective”. In:
- Khalidi, Muhammad Ali (2013). *Natural Categories and Human Kinds : Classification in the Natural and Social Sciences*. New York: Cambridge University Press.
- 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.
- 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.
- Ney, Alyssa (2021). “Reductionism”. In: *The Internet Encyclopedia of Philosophy*. URL: <https://iep.utm.edu/red-ism/>.

- 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.
- 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 (2010). “Monism: The priority of the whole”. In: *The Philosophical Review* 119.1, pp. 31–76.
- (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.
- 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.
- 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.

- 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.
- 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. URL: <https://doi.org/10.1007/s11229-018-01998-6>.