

Theoretical Relicts: Progress, Reduction, and Autonomy

Katie Robertson and Alastair Wilson

When once-successful physical theories are abandoned, common wisdom has it that their characteristic theoretical entities are abandoned with them: examples include phlogiston, light rays, Newtonian forces, Euclidean space. But sometimes a theory sees ongoing use, despite being superseded. What should scientific realists say about the characteristic entities of the theories in such cases? The standard answer is that these ‘theoretical relicts’ are merely useful fictions. In this paper we offer a different answer. We start by distinguishing horizontal reduction (in which a superseded theory approximates the successor theory) from vertical reduction (in which a higher-level theory abstracts away from the lower-level theory, but nonetheless can be constructed from it); these are usually regarded as having different ontological consequences. We describe a ‘verticalization’ procedure that transforms horizontal reductions into vertical reductions. The resulting verticalized theories are abstractions rather than approximations, with restricted domains. We identify a sense in which the higher-level theory describes distinct subject matters from the lower-level theory, enabling in certain cases the higher-level theory to retain distinctive explanatory power even in the presence of reduction. We suggest that theoretical entities from superseded theories should be retained in a scientific realist worldview just when, reinterpreted as higher-level abstractions, those theories and their characteristic entities continue to perform distinctive explanatory work in providing the best explanation for less fundamental phenomena of interest. In slogan form: a good relict is an emergent relict.

ORCIDs: Robertson, <https://orcid.org/0000-0002-0645-0761>; Wilson, <https://orcid.org/0000-0003-3027-5016>

1. Introduction

Physics makes progress: this much is uncontroversial. As part of this progress, our view of the physical world expands: we have discovered atoms, nuclei, quarks, galaxies, pulsars, black holes. It sounds equally uncontroversial to say that sometimes, as part of progress, our view of the physical world contracts. When once-successful physical theories are superseded, common wisdom has it that their key theoretical entities are abandoned. Classic examples are caloric, phlogiston, and Aristotelian natural motions of bodies.

Sometimes, though, entities are retained in scientific descriptions despite being relegated to a less fundamental role. We still describe and explain the world in terms of heat (as opposed to kinetic energy) when we use thermodynamics, rays of light (as opposed to photons) when we use geometrical optics, and space and time (as opposed to spacetime) when we launch satellites and fly aeroplanes. A theoretical relict is a theoretical entity posited by a superseded yet once-successful scientific theory. Some theoretical relicts continue to play some role in our scientific practice but are standardly thought to be useful fictions; others are wholly abandoned. Under what circumstances should scientific realists acknowledge the existence of a theoretical relict corresponding to an entity posited by some previous candidate for a more fundamental theory? Why, for example, should scientific realists maintain that there are atoms and light rays but that there are no vital forces and no phlogiston? We call this the ‘puzzle of theoretical relicts’.

In addressing the puzzle, we aim to set aside issues about reference of theoretical terms insofar as is possible. Our question is not when we have continuity of reference across theory change: it doesn't matter for our purposes whether the same term actually gets used in the same way before or after a theoretical change. Continuity of reference can be seemingly too easily secured given a broadly causal theory of reference, and the question has led to plenty of mostly inconclusive discussion. Hardin and Rosenberg ([1982]) argue that ‘ether’ refers to whatever causes the phenomena of electromagnetic radiation; Bird ([1998], [2000]) worries that a causal approach risks ‘phlogiston’ ending up referring to oxygen, while Kitcher ([1993]) embraces that outcome; Myrvold ([2020]) offers a nuanced discussion of the semantic questions involved in the case of the luminiferous ether and caloric. But we think, following Ladyman ([2011]), that successful reference is largely beside the point for explaining why an old theory was successful. A better explanation cites shared structure between the theories as demonstrated by correspondence principles (Post [1971]) and—most directly—by inter-theoretic reduction.

We begin in section 2 by distinguishing between horizontal and vertical reduction. Central to all putative reductions is the construction or derivation of one theory from another. In horizontal reduction, an improved successor theory explains the success of a predecessor theory; in vertical reduction, a higher-level theory is derived or constructed from a lower-level theory. These two types of reduction have been the focus of two largely distinct bodies of literature; horizontal reduction is closely associated with the theory change debate, whereas vertical reduction is central to debates over physicalism.

With our distinction between types of reduction in hand, in section 3 we explain a sense in which a vertically reduced theory has a distinct subject matter from the underlying theory, and in section 4 we describe a procedure of verticalization that converts horizontal reductions to vertical reductions. Verticalization permits a reconceptualization of the old, approximate theory as a higher-level, more abstract, theory. The last part of section 4 considers when and why we should deploy this procedure, by reference to the explanatory role of the verticalized theory.

Section 5 shifts focus from the verticalized theory to the entities characteristic of that theory. The received view of horizontal reduction is that scientific realists typically discard an old theory's characteristic entities when that theory is superseded. But the received view of vertical reduction is that it needn't lead to elimination; instead most contemporary physicalists (*contra* Churchland [1981]; Churchland [1986]) are at pains to retain higher-level entities as part of a scientific realist ontology. What implications, then, does verticalization have for the ontology associated with the reduced theory? We suggest that scientific realists should retain theoretical entities from horizontally reduced theories when, reinterpreted as higher-level abstractions after verticalization, the entities can perform distinctive explanatory work in providing explanations for phenomena of interest. Section 6 addresses an objection recently raised by Saatsi ([2022]) to 'effective' forms of scientific realism like our own. We argue that our approach avoids reifying the theoretical inconsistencies against which Saatsi warns; verticalization's starting point of successful horizontal reduction is crucial here. Section 7 is a conclusion.

The puzzle of theoretical relicts is ultimately resolved, not by identifying any universal metaphysical feature that relicts have in common, but by first harmonizing the representational roles of different scientific theories and their characteristic entities, and then identifying a distinctive explanatory benefit that can justify their retention. This explanatory benefit derives not from the status of relicts as obsolete candidates for our most fundamental theory of some phenomena, but from their status as successful higher-level explainers of aspects of those phenomena. In this respect our proposal has links to recent work on emergence that uses criteria of explanatory autonomy and novelty to identify candidates for inclusion in the scientific realist ontology. This permits a slogan formulation of the proposal: a good relict is an emergent relict.

2. Two Concepts of Theory Reduction

Some old scientific theories, especially those that pre-date modern science, are straightforwardly abandoned; they are never used in practice and their characteristic theoretical entities are discarded. The four humours theory of the human body, which placed yellow bile and black bile alongside blood and phlegm as basic explanatory elements, turned out to be nothing more than a dead end. So did Aristotelian physics, with its characteristic posits of a body's natural place and corresponding natural motions. But many superseded theories enjoy substantial empirical success, and aspects of their descriptions of the world are retained in subsequent theorizing (Hoefer and Martí [2020]). Classical electrodynamics was good enough to help us design and build radio sets and fax machines. Newtonian mechanics enabled startlingly accurate predictions of astronomical phenomena like comets and is still used to get rockets to the moon. How could these theories have been—and remained—so successful in spite of being false? A familiar answer appeals to inter-theoretic reduction of the horizontal kind.

Here we take reduction to be a relation that in the first instance holds between two theories, T_t and T_b . This choice amounts to taking a 'theory-first' approach to reduction (Batterman [2001]; van Gulick [2001]), rather than an ontology-first view (for example, Oppenheim and Putnam [1958]). Different accounts of reduction further specify the nature of the reduction relation: as definitional extension (Nagel [1961]; Butterfield [2011a], [2011b]) or translation (Dewar [2019]), as a structural relation (Suppes [1967]), as a local, *a posteriori* relation (Rosaler [2015], [2017]), or as an analogical relation (Bickle [1998]).¹ The key feature common to all major accounts of reduction is that the reductive relationship in some sense permits recovery of some elements of T_t from some elements of T_b . 'Recovery' is a particularly apt metaphor here since T_t can guide our search for how to get from T_b to T_t . This is particularly relevant when limits are crucial to the reduction.

The kind of reduction we have in mind can be flexible in various ways. We discuss shortly how only an approximate 'cousin' of the original theory need be found (Schaffner [1967]). Another flexibility is that reduction is typically local (Rosaler [2015], [2017]): one particular description of a certain type of system may be reduced to another particular

¹ Part of the dispute about reduction is how distinct these proposals really are (Dizadji-Bahmani et al. [2010]; van Riel and van Gulick [2019]).

description, rather than grand swathes of a theory being reduced in one fell swoop. We need not draw a strict line between whole theories, models, or fragments of a theory (Crowther [2018])—indeed, the idea that reduction is local goes back at least to Lewis ([1980]). A key development proposed by Schaffner ([1967]), and endorsed by many (Butterfield [2011a], [2011b]; Dizadji-Bahmani et al. [2010]), is that the reduction relation can also be flexible in that only an approximate cousin of T_t need be found, or recovered, starting from T_b . As such, the details of a given reduction might show how to correct T_t ; one important consequence of reduction is that it shows us how T_t and T_b are, despite first appearances, compatible with each other.

For the purposes of this paper, we won't need the fine-grained details of these accounts of the reduction relation. The crucial feature of reduction in our sense is that the dynamical laws, or what we might more generally call the 'nomological structure' of the old theory, are recovered. To reflect this, we operate with a flexible notion of reduction as construction (Robertson [2019]):

Reduction as Construction

Theory T_t is reduced to T_b if the equations, quantities, and variables of T_t can be constructed from the equations, quantities, and variables of T_b .

The usefulness of the construction approach to reduction is that it makes explicit that theory reduction may include adding in additional assumptions of different types alongside the 'core' equations of each theory. It acknowledges that it is difficult to secure reductions; examples are few and far between, and so whatever resources are required—approximation, idealizations, mathematical procedures like limits, and novel conceptual resources like probabilities and initial conditions—are sanctioned.² Much of the debate about particular reductions stems from whether the additional resources are acceptable or not; for example, the debate on whether the use of limits blocks reduction (Batterman [1995], [2001]; Butterfield [2011a], [2011b]).

Our invocation of reduction-as-construction places the bar for successful reduction comparatively low, lower than some would be willing to accept. For example, Frigg and Werndl ([2019]) maintain that the assumptions we add in must be 'physical' in some

² For a response to the objection that this trivializes reduction, see (Uffink [1996]).

suitable sense. But for present purposes, the low standard for reduction is dialectically innocent, since it brings more candidate reductions into the scope of our argument; if we can establish our conclusions in the context of a flexible approach to reduction like this one, then they ought to carry straight over to more demanding conceptions of reduction.

Contrast two types of reduction fitting our schema:

Horizontal Reduction

T_i is an old, tainted theory that is reduced to a newer, better theory, T_b . ($T_b =$ 'better'.)

Vertical Reduction

T_i is a macroscopic, higher-level, or top theory that is reduced to an underlying, microscopic, or bottom-level theory, T_b . ($T_b =$ 'bottom'.)³

If the superseded theory can be approximately recovered—that is, constructed—from the successor theory under suitable conditions, then the old theory has been horizontally reduced. A successful horizontal reduction can account for the empirical success of the old theory. In such reductions, the laws/regularities of the old theory approximate those of the newer theory, and this approximation enables us to understand why the old theory enjoyed the success it did.

In a classic type of case, we have the old theory recovered as a limiting case of the new theory—for example, classical mechanics is recovered from special relativity in the low-relative-velocity limit. But not all cases of horizontal reduction need take a limiting-case form: for example, one might horizontally reduce the macroscopic description of a game of life scenario (Gardner [1970]) to the microscopic description: macroscopic rules governing the behaviour of gliders will turn out to be approximately recovered by the micro-dynamics of the game provided that the appropriate initial conditions are included in the construction. Imagine a community that first studies the macroscopic behaviour of the game of life, before starting to make inferences about what underlies it and discovering the underlying deterministic dynamics.

³ Crowther ([2018]) employs alternative terminology directed at a similar distinction, distinguishing diachronic (our horizontal) reduction versus synchronic (our vertical) reduction. Franklin-Hall ([2016]) also uses a vertical–horizontal distinction.

An example already discussed is classical mechanics and special relativity in the low-velocity limit. We can construct the equations of motion of classical mechanics by taking the limit of v (the velocity of an object in our frame of reference) to be low relative to c (the speed of light); the v^2/c^2 term in the special relativistic equations becomes negligible in the low v/c limit. For example, in the low velocity limit, the relativistic law for velocity addition,

$$u = \frac{u' + v}{1 + \frac{u'v}{c^2}}$$

becomes the Newtonian velocity addition rule: $u = u' + v$. Sometimes physicists speak loosely about c tending to infinity; but of course, the speed of light doesn't vary. Rather v/c tends to zero: the velocities of systems like footballs (or even rockets) are so small compared to c that they are effectively zero. The interpretation usually given to this limit is that the classical description approximates the more accurate special relativistic description: the approximation is excellent for low relative velocities (as with billiard balls and artillery shells) but becomes extremely poor once v/c is not negligible, and at relative velocities close to c (as in the case of cosmic rays) fails completely.

In horizontal reduction, classical mechanics and special relativity have the same subject matter in a natural sense, one that we aim to make more precise in the next section. The new theory accounts for the relevant aspects of the world—what we will call the target phenomena—in a way that strictly improves on the account of those aspects offered by the old theory. One way of emphasizing that the two theories have the same target phenomena is that these theories are naturally seen as competitors. It is this comparison between the descriptions of the same target phenomena that enables us to establish that T_b is more accurate than T_t , but that nonetheless the descriptions given by T_t approximate those given by T_b . In short, horizontal reductions hold between theories that describe the same target phenomena, but which use varying degrees of approximation and hence achieve varying degrees of accuracy. Once the dust has settled on how the horizontal reduction proceeds, the scientific realist often wants to say that while they are committed to the new theory and its characteristic entities, the old theory and its entities can be relegated to the status of mere useful fictions. (We discuss our preferred alternative to this approach in section 5.)

In contrast, approximation is not part of the vertical reduction relation. Is hydrodynamics less accurate than particle mechanics? Is biology less accurate than chemistry? No; and these questions strike the listener as misguided. In order to compare theories with respect to accuracy, the theories need to make predictions about the same phenomena. But in vertical reductions the descriptions given by T_t are about a different subject matter than the descriptions given by T_b , and so there is no straightforward sense in which one theory is more accurate than the other. Rather, the theories characterize different aspects of reality; to take an example we will consider in depth, statistical mechanics characterizes the bulk properties of matter whereas the underlying classical dynamics tracks the position and momentum of each molecule in the gas. Or, to take another example, hydrodynamics is concerned with the bulk motion of a fluid, whereas particle mechanics is concerned with the individual components of a fluid. As such, the higher-level theory T_t description is more abstract, describing reality in less specific and more general terms. Such descriptions might be more suitable for a variety of purposes: practical, experimental, theoretical.

The relationship between classical mechanics and statistical mechanics is a paradigmatic example of vertical reduction. Statistical mechanics abstracts away from some of the microscopic details—such as the position of molecule #11067395—and instead focuses on the macro-patterns, such as how quickly a gas will spread out throughout a container. As the higher-level theory T_t , statistical mechanics can be derived or constructed starting from the underlying micro-dynamics of classical mechanics. In this case study, the irreversible higher-level equations can be constructed from the underlying micro-dynamics. This involves coarse-graining—a form of abstraction—of the description of the system, ρ . Certain assumptions are required, but we can construct, or derive, an irreversible equation such as the Boltzmann equation, which describes a gas relaxing to equilibrium. An initial condition is required, and it is hard to find an abstraction—a coarse-graining of the state—that leads to genuine macro-dynamics. There are many bad

$$\begin{array}{ccc} \rho_{cg}(t_0) & \xrightarrow{C} & \rho_{cg}(t) \\ \uparrow \hat{P} & & \uparrow \hat{P} \\ \rho(t_0) & \xrightarrow{U} & \rho(t) \end{array}$$

Figure 1. ρ is the full probability distribution at different times; ρ_{cg} is the coarse-grained probability distribution, U is the micro-dynamics, C is the macro-dynamics and \hat{P} is the coarse-graining map.

choices. For successful abstraction in the case of statistical mechanics dynamics, we need the higher-level macro-dynamics to ‘mesh’ in a certain way with the micro-dynamics—so that the micro-details omitted by coarse-graining are not dynamically relevant. If this meshing condition holds, then the lower-level details truly do not matter.

However, we know that there are limitations to the situations where this meshing condition holds. In particular, due to Poincaré’s recurrence result, we know that given long enough any classical or quantum system will return to a state arbitrarily close to its earlier state: the earlier state will ‘recur’. At the recurrence time, the irreversible equations of statistical mechanics are no longer apt: the system has returned to its earlier low entropy state, but the macro-dynamics decree that entropy can only increase. So we know the meshing situation breaks down at the recurrence time, and this is because one of the assumptions made in the construction, the Markovian approximation (Zeh [2007]; Wallace [2012], [unpublished]; Robertson [2020]), no longer holds. The other assumption required is an initial state condition: at t_0 , the probability distribution ρ must be suitably well behaved, otherwise we cannot find autonomous dynamics.⁴ Whether there is a successful vertical reductive relation between these theories hinges on how well justified these two assumptions are, and (for a Nagelian reduction) on whether they qualify as bridge laws in the appropriate sense.

⁴ For the dynamics to be autonomous, they must not functionally depend on the ‘irrelevant’ variable; see (Robertson [2020]).

Katie Robertson and Alastair Wilson




Figure 2. Emergent chance: deterministic lower-level histories are compatible with higher-level indeterminism. Reproduced from (List [2014]).

This case study of classical mechanics and statistical mechanics is an example of a more general schema, explored by List ([2019]), in which different theories can describe and explain phenomena at different levels by asking and answering different questions.⁵ The descriptions applicable at different levels might be very different in character. For example, if we abstract to a coarser description, the dynamics might be probabilistic rather than deterministic—or vice versa. These possibilities are illustrated in figure 2.

In the case of the classical mechanics–statistical mechanics reduction, the micro-dynamics are time-reversible, but the macro-dynamics of statistical mechanics are irreversible (Robertson [2020]). By identifying the right abstraction maps between state spaces, the relationships between macro- and micro- (or higher-level and lower-level) states are rendered especially clear; the different levels neither compete nor conflict. It remains an open question what the ontological consequences of such a vertical reduction are. We think that it is now a minority view that the higher-level ontology is eliminated.⁶ In any case, the approach we will be exploring is non-eliminative about higher-level entities.

⁵ In addition to the non-equilibrium statistical mechanics case, we can consider the reduction of equilibrium statistical mechanics to classical mechanics along the same lines. But, for example, some hold that the adding of probability is too suspicious to be considered as a bridge law (Sklar [1993]; for a rebuttal, see Uffink [1996]), or others claim that the initial conditions stipulated as part of the reduction merely shift ‘the lump in the rug’ (Price [1996]). We set such concerns aside here.

⁶ For further discussion, see (Ney [2008]). Contemporary eliminativists include French ([2014]), Esfeld and Deckert ([2017]), and Carroll ([2021]).

3. Reduction and Subject Matters

In the previous section we characterized approximation as key to horizontal reduction, and abstraction as key to vertical reduction, generating the following contrast:

Horizontal Reduction

The older theory approximates the newer theory. It provides a less accurate, but still successful, description of some common subject matter.

Vertical Reduction

The higher-level theory abstracts away from the more detailed lower-level theory. It provides successful descriptions of a different, higher-level, subject matter.

In contrast to horizontal reductions, two theories in a vertical reduction are never seen as competitors. The higher-level or macroscopic, description complements the lower-level description, rather than competing with it.

What is a theory's subject matter? We understand subject matter in purely modal terms, along the lines indicated by Lewis ([1988a], [1988b]). Subject matters are, roughly, ways of partitioning possibilities so that each different way for the subject matter to be corresponds to a single cell of the partition. There is a close connection between this Lewisian notion of subject matter and the logic of questions and answers: a question is associated with a partition corresponding to all of its possible answers (Lewis [1988b], p. 162). While it does not draw hyperintensional distinctions, as do many accounts of subject matter (Hawke [2018]), this core connection to questions will be sufficient for our purposes in this paper.⁷ It accommodates both variation in grain and variation in particular domain (how things are biologically in Wales, or how things are economically in mature capitalist economies).

The modal account of subject matters provides a basic link between subject matters and distinctions between possibilities, but it does not tell us which possibilities to attend to in the case of any specific theory. Which partitions of possibilities correspond to the

⁷ Indeed, overly fine-grained accounts of subject matter will return the result that every scientific theory—perhaps every individual application of a scientific theory—has a distinct subject matter from every other.

subject matter of a particular scientific theory on this model? To answer this question, we turn to the notion of a theory's 'target phenomena'—roughly, whatever causes the observational data that are used to test the theory (cf. Bogen and Woodward [1988]). Examples of target phenomena are intended to be familiar and intuitive. Newtonian mechanics aimed to describe the same phenomena as special relativity: the relative motion of material bodies. Phlogiston theory aimed to describe the same phenomena as oxygen theory: combustion and suffocation. The ray theory of light aimed to describe the same phenomena as the wave theory of light: rainbows, prisms, lenses.

With the notion of target phenomena in hand, a theory's subject matter for our purposes can be identified with the distinctions it draws amongst possibilities concerning the target phenomena; or, equivalently, with the questions that it allows us to ask and to answer about the target phenomena. At what temperature does the kettle boil? Why does the flame turn green when this barium compound is tested? Why did the football's trajectory take it through the window? These questions each define (subject to the usual vagueness of descriptive terms) a partition over possible worlds; each cell corresponds to one possible answer. The subject matter of a theory is how things stand with respect to the set of questions connected to the target phenomena; each cell of the combined partition corresponds to a possible combination of answers to these questions.

We can now transpose our distinction between horizontal and vertical reduction to the present framework. The questions that an old theory aims to answer are the same as for the new theory: they are generated by the same target phenomena. Higher-level theories and lower-level theories, by contrast, aim at different sets of questions: they have different target phenomena. In the latter case, the set of questions answered by the higher-level theory T_t overlaps with T_b 's set. This is to be expected: their subject matters are not orthogonal (compare: the number of universities, the number of jellyfish) since there are modal connections (such as supervenience) between different levels.

A key feature of our conception of vertical reduction is that higher-level theories have distinct (not necessarily disjoint) subject matters from the lower-level theories that feature in the reduction of the higher-level theories. This thesis of distinctness of subject matter comes in two forms, one controversial and one uncontroversial. The controversial form of the thesis says that the higher-level theory is about phenomena that the lower-level theory is not about: there is new content in the higher-level theory not present (even implicitly) in the lower. The uncontroversial form of the thesis, which is all that we will

need for our main argument, says that the lower-level theory is about phenomena that the higher-level theory is not about. This distinctness is assured by the restriction step of the verticalization procedure described in the next section.

We aim to remain neutral on the controversial form of the distinctness thesis, which is associated with more robust forms of anti-reductionism. One such approach is Cartwrightian pluralism (Cartwright [1999]), according to which higher-level theories describe phenomena that cannot be described at all using lower-level theories. In the case of the reduction of statistical mechanics, this amounts to the view that statistical mechanics describes phenomena that quantum mechanics can't describe (Hartmann [2000]). But how should we understand a claim like 'quantum mechanics describes gases'? Should it be understood as 'gases are in the domain of applicability of quantum mechanics' or 'quantum mechanics provides answers to questions about the rate of relaxation to equilibrium'? The former looks plausible; the latter looks implausible. Quantum mechanics applies to a gas, and with the help of additional assumptions we can construct the Boltzmann equation from quantum mechanics. This was the vertical reduction discussed above. But what do we do if we want an answer to the question: how quickly does this gas reach equilibrium? We could either just use the Boltzmann equation (that is, take the original statistical mechanics answer to the question), or we could attempt to evolve the full probability distribution according to the micro-dynamics and then coarse-grain (see fig. 1).⁸ The former strategy is clearly viable, and the latter strategy clearly unviable; accordingly, it is undeniable that statistical mechanics does a much better job at describing gases relaxing to equilibrium. At best, quantum mechanics describes this phenomenon only with help from statistical mechanics.⁹

The upshot is that the descriptions a theory might offer in practice are an impoverished subset of the descriptions/answers it provides 'in principle'. This impoverishment might be due to human factors; some equations are psychologically difficult for us to solve. But there might also be deeper explanations of the divergence between principles and practice. For example, solving the Schrödinger equation for a typical gas that contains 10^{23}

⁸ Whether relying on 'coarse-graining' implies that the higher-concepts (and theory) are implicated in the answer is a vexed question, but at least it is clear that there are not two equations (one translated from the other) that we compare (and can thus see which is more accurate).

⁹ The concept of a macro-state in the classical mechanics–statistical mechanics reduction is a particularly clear case in which higher-level concepts apparently need to be imported by hand in the construction of the higher-level theory in terms of the lower-level theory.

molecules is not just a mathematician's headache; it is hard to see how it could be computed given what a vast number 10^{23} is, considerably larger than the number of grains of sand on all the beaches on Earth (roughly 10^{18}). Yet the idea that the higher-level theories do better at answering certain questions than the lower-level theories to which we reduce them is still compatible with the uncontroversial distinctness thesis. The uncontroversial distinctness thesis allows us to set aside these difficult issues about whether quantum mechanics really describes a gas, or a cell. Instead, our modal account of subject matters can distinguish the subject matters of T_t and T_b even if we work with a very liberal 'in-principle' sense in which T_b describes the same phenomena as T_t —all that is required is that there is a subvening description in T_b .

The uncontroversial form of the distinctness thesis says that the lower-level theory describes phenomena that cannot be described at all using the higher-level theory: the content of the micro-physical description outruns the content of the higher-level description. Chemical reactions cannot be described at all in terms of evolutionary biology. In successful abstraction, the higher-level description manages to capture a different (and explanatorily useful) set of distinctions between possibilities precisely in virtue of tactically leaving out some of the detail of the micro-physical description. This inevitably results in certain physical facts that have a description at the lower level not getting described at all at the higher level: in the case of reduction of statistical mechanics to quantum mechanics, these facts include what happens at the recurrence time, what happens with unusual initial conditions, and more generally what happens to systems that do not fulfil the conditions of applicability of the statistical mechanics description.

Quantum mechanics answers more questions than statistical mechanics, and this is the uncontroversial sense in which the two theories have distinct subject matters. But even in the presence of a successful vertical reduction the statistical mechanics description continues to give better answers than the underlying description does to certain questions. How quickly will this gas reach equilibrium? What is the highest power output this steam engine can achieve? What's the lowest temperature that can be reached by evaporation in current humidity levels? If we want to answer the question 'how quickly does this sample of gas relax to equilibrium?' we use the Boltzmann equation. As discussed above, there is no alternative equation available (no 'microscopic translation' of the Boltzmann equation) into which we can feed micro-physical information about the gas particles,

other than to use the micro-dynamics and then coarse-grain. The abstraction leads to distinctive explanatory power by uncovering the macro-dynamics.

The set of questions for which the higher-level theory gives the best explanations has sometimes been termed the theory's 'proprietary explananda' (Woodward [2021]). This is connected to the idea that higher-level captures less; statistical mechanics abstracts away from the full micro-physical details, in particular by using probability distributions, and not tracking some of the details (for example by averaging, or as we saw earlier, by coarse-graining). It is partly by purposefully neglecting those questions that statistical mechanics succeeds so well in describing and explaining its proprietary subject matter.

Let us take stock. So far, we have:

- (1) Distinguished two types of reduction, horizontal and vertical—the former between competing theories about some common target phenomenon, the latter between complementary theories that explain that phenomenon at different levels of detail.
- (2) Characterized the relation between vertically reducing theories in terms of their having distinct subject matters, according to a simple modal framework.
- (3) Drawn an initial link between the variation in subject matters of different vertically related theories and the variation in explanatory questions raised by the target phenomena of those theories.

This puts us in a position to characterize in the next section a general recipe for generating vertical reductions out of successful horizontal reductions.

4. Verticalization

Horizontal reductions are distinct from vertical reductions, but the two have a close relationship. We will argue that horizontal reductions can often be transformed into vertical reductions, through a procedure that we call verticalization. In our terminology, 'verticalization' in the first instance applies to reduction relations between theories. We feed in a horizontal reduction relation between an old and a new theory and we get out a vertical reduction relation between a higher-level and a lower-level theory. But it is also natural to talk of the 'verticalized theory'.

How does verticalization work? There are two key components to the verticalization procedure; we re-invent the old theory by restricting it and reinterpreting it. First, a restriction is imposed on the domain of the old theory: the verticalized theory is assigned a content that is wholly about some restricted range of phenomena. The domain of the restriction is typically somewhat vague: to systems with small v/c , or to systems with actions large relative to Planck's constant. While at the level of the theory, it is hard to give a precise answer about the domain (since *inter alia* there is vagueness in 'small' and 'large'), the picture looks more precise in individual cases, that is, when we have reductions between particular models or descriptions of particular systems; then we can have a clear handle on when one model is a good approximation of another (Rosaler [2015]; Wallace [unpublished-c]). And as in the case of effective field theories, we are interested in the cases far from the boundary (or 'cut off' in the EFT language)—individual models will be far from this regime.

Second, the descriptive apparatus of the old theory is reinterpreted as a higher-level, coarser-grained, more abstract description. As part of this, we limit the degree of precision with which the question is answered. In this reinterpretation process, properties regarded by the old theory as (relatively) fundamental might now be regarded as bulk properties of something deeper. Some properties from the old theory might be dropped altogether in the move to the verticalized theory, if they prove to be nothing more than an artefact of the older model.

The restriction step means that the old (soon to be higher-level) theory answers fewer questions than the new/lower-level theory. This step is necessary to respect the uncontroversial kind of lack of overlap of subject matter noted in the previous section: isolated non-chaotic single-particle systems, described easily by quantum mechanics, are simply beyond the scope of statistical mechanics. We accordingly need to restrict the domain of the old theory so as to limit it to the circumstances in which it gets some nomological structure right, and so latches onto genuine dependencies. This is straightforward to implement within the modal account of subject matter. The partition giving the subject matter of the verticalized theory is indifferent to the behaviour of systems outside the theory's scope, with all worlds that differ only in respect of the behaviour of those systems being placed in the same cell of the relevant partition.

The reinterpretation step means not only that the partition is cast over a smaller subset but also that the partition is coarser-grained, drawing fewer distinctions amongst possibilities. This step is necessary to make sense of an ongoing explanatory role of the verticalized theory. Understood as a failed attempt to correctly describe the phenomena, there is no apparent reason why a high-level theory should be successful. But understood as a successful attempt to correctly describe more abstract aspects of the phenomena—contained within its newly limited domain—the path is open to vindicating a continued role for the verticalized theory, and its entities, in the scientific enterprise. Again, the modal approach to subject matters makes this simple to implement. Instead of regarding the theory as specifying (via the questions associated with its target phenomena) the finest-grained partition of worlds, we now regard it as specifying a coarser-grained partition—a more abstract description—which places worlds that agree with respect to the macro-description into the same cell, even though these worlds disagree with respect to the micro-description.

Together, the restriction and reinterpretation steps give T_i a distinct subject matter from T_b . Now T_i is understood as an ‘effective theory’ (in the sense of an effective field theory; cf. Williams [2019]). A characteristic feature of an effective theory is that it aspires to describe phenomena only within a limited range and to a limited range of accuracy; and moreover the effective field theory explicitly specifies the scales on which it breaks down. As Williams ([2019], p. 222) puts it: ‘EFTs thus provide formal signposts delineating the physical domains in which one should and should not trust the theory to provide reliable ontological guidance’.

We can use an example from List ([2017]) to illustrate the contrast between approximation and abstraction. List’s comprehensive and precise account of levels is an expansion of the Lewisian account of subject matters we draw upon here, and the framework centrally includes abstraction mappings between levels. He draws on an example from decision theory: the level of awareness associated with an agent depends on the fineness of the distinction the agent is able to draw. ‘The agent is aware of some feature of the world (or a feature of some item) if and only if he or she is able to distinguish worlds (or items) with that feature from ones without it’ (List [2017], p. 8). Greater awareness is reflected in a finer-grained partition over possible worlds. Crucially, instead of thinking that the less aware agent merely approximates the more aware agent, in decision theory

we can conceive this as abstracting away from the differences that the more aware agent can discern. Approximation is reinterpreted as abstraction.

An objection looms large. Can't we always vindicate any old theory by reinterpreting it, to understand it as applying only within certain domains and/or to certain degrees of accuracy? After all, it sounds like a truism that all theories are successful insofar as they get anything right about anything at all: a theory with no empirical success at all would not make it off the blackboard. The answer is yes; in principle, we can verticalize the Thalesian theory that everything is made of water to restrict it to apply only to ice cubes, and we can verticalize the phlogiston theory of combustion by (for example) reinterpreting phlogiston to identify it with an absence of oxygen. But in our view the in-principle-availability of the verticalization procedure does not open the floodgates to just any old verticalized theories and entities: there are restrictions on when verticalization is appropriate.

When should a horizontal reduction be verticalized? One naïve suggestion would be: whenever the old theory got anything right at all. But that is a risky approach to take for the scientific realist. After all, phlogiston theory got some nomological structure right: indeed, Noretta Koertge ([1969]) argues that phlogiston theory can be (horizontally) reduced to redox theory (see also Ladyman [2011]). So: does phlogiston exist after all? If not, what is the difference between this case and cases where we do think the reduction should be verticalized—such as the reduction of thermodynamics to statistical mechanics or classical mechanics to special relativity? The difference, we think, is an explanatory one.

Our guiding thought is that some horizontal reductions—but not all!—show the old theory to be explanatorily redundant once we have the new theory in hand. For instance, there are no relevant explananda for which phlogiston theory does better than redox theory, but there are some relevant explananda for which Newtonian mechanics does better than special relativity. (Why did my ball miss the goal? Is this pendulum's motion chaotic?) Our suggestion, then, is that a horizontal reduction should be verticalized just when the verticalized theory T_t gives better explanations than the lower-level T_b for some rel-

evant explananda, that is, when T_b does not strictly dominate T_t with respect to explanatory power.¹⁰ We understand this as an empirical constraint, and in principle revisable: if we later discover some valuable new explanation that the verticalized theory could offer us, we should reconsider whether that reduction should be verticalized after all.

Some restriction of relevance on explanatory questions is required. Our approach breaks down if all potential explananda count as relevant, including purely theoretical ones from the old theory. How much phlogiston is there in this air? Is the Earth at absolute rest? These questions seem to have no empirical basis since they were generated not by the phenomena, but only by a theory that was later abandoned. To account for relevance, we propose to draw on the notion of a target phenomenon discussed in section 3; an explanandum is relevant if it is part of a target phenomenon of scientific interest.

It may help to distinguish two cases where an old theory will fail to have a verticalized higher-level correlate. One way for an old theory to fail to verticalize is for there simply to be no successful horizontal reduction in the first place. In the case of radically defective old theories, the correct response is abandonment rather than rehabilitation; an example is the Titius–Bode ‘law’.¹¹ But a different and more interesting way to fail to verticalize is for there to exist a horizontal reduction, but for there to be no distinctive explanatory power of the reduced theory. Here phlogiston is again our example. Although one can identify correspondence principles between phlogiston theory and successor theories, the phlogiston theory no longer provides the best explanation of any phenomena of (independent) interest. By contrast, the reduction of thermodynamics to statistical mechanics for us is a paradigm case of the vindication of the old theory through judicious verticalization: thermodynamics is frequently the best level of explanation and description for practical scientific and engineering purposes.

¹⁰ *Contra* Frigg and Werndl ([2019]), we assume that effective theories can explain. We also set aside the potential concern that the existence of Kuhn losses means that the ‘does not strictly dominate’ condition is always satisfied; we assume the realist can account for whatever Kuhn losses there may be in terms of some contingent features of the scientific process. For discussion, see (Hartmann [2000]).

¹¹ The relevant theory here is one that identifies a coincidence with respect to the spacing of planets in the solar system and mistakenly elevates this coincidence to the status of a law. The regularity in actual planetary spacings cited by the ‘law’ is in some sense explained by our best current causal-historical story about solar system formation. However, we no longer regard this regularity as lawlike: there are no genuine dependencies that the bad planetary generalization latched onto in the first place. Accordingly, we don’t have a case of reduction between theories here.

Note there's one assumption that this verticalization criterion assumes: that it is possible that sometimes the better explanation resides at the higher level. If the best explanation were always at the fundamental level (Railton [1978], [1981]), then the motivation to take the higher-level ontologically seriously—as we will do in the next section—would be greatly reduced. That higher-level explanations are sometimes better than lower-level explanations has been widely endorsed in the philosophy of science since at least Putnam ([1975])—see also (Kitcher [1981], [1984]; Batterman [2001]; Strevens [2008])—but what makes higher-level explanations better remains controversial. For us, the problem becomes: what do we mean by 'greater explanatory power' in our verticalization criterion? There are various familiar respects in which one explanation might be better than another, and some of them are rather parochial and contingent. For scientific realist purposes, what is needed here is an objective criterion of explanatory betterness—albeit one which might still include some user-friendliness elements such as calculational tractability, simplicity and the like—which in some circumstances identifies higher-level explanations as better than lower-level explanations. We will now briefly canvass some potential criteria of this kind to give an idea of what we have in mind, but ultimately we aim to remain neutral on exactly what accounts for the value of the explanations provided by higher-level theories.

'Proportionality' is a widely discussed criterion, first discussed by Yablo [1992], and based around the notion of difference-making. In at least some circumstances it counts high-level explainers as better, and it is not tied to subjective interests of explainers beyond their initial choice of contrastive explanatory question.¹² The particular account of proportionality developed by Menzies and List ([2009], [2010]) in their response to Kim's ([1998]) causal exclusion problem is a simple approach that could be employed to account for the distinctive explanatory power of high-level explanation; a related approach is described by Woodward ([2008], [2021]). The presence of oxygen in the room is part of the most proportionate explanation for why the metal tarnishes; the exact configuration of quantum fields is not. However, the proportionality approach has competi-

¹² Weslake ([2010]) and Franklin-Hall ([2016]) object to Woodward's ([2008]) application of proportionality, offering counterexamples based around models with non-standard variable choices. Woodward ([2021]) replies by acknowledging the importance of the right choice of variables, and emphasizing the need for the model to be apt in that it correctly represents the dependence structure out there in the world.

tors, and our project could equally be combined with related approaches such as Strevens's ([2008]) kairetic account of causal explanation and its minimality condition. For Strevens, the best explanation is the most abstract that still successfully explains the explanans. For our broader purposes in this paper, we do not need to commit to any particular one of these accounts of the source of the explanatory value of higher-level descriptions: they will have the same upshot for central cases from the literature, by vindicating the explanatory value of atoms and thermodynamic properties without vindicating the explanatory value of phlogiston or caloric.

What is the role of the verticalization procedure we have described? Our motivation, at least, has been to understand scientific explanatory practice: in particular, to account for the ongoing explanatory utility of surpassed scientific theories. We take it that explanations have to be true (at least in their core elements) to qualify as genuine explanations. If Newtonian mechanics is flatly false, it can't explain; and yet it does explain. A natural response is to suggest that Newtonian mechanics is 'approximately' true; but some worry that these lingering elements of falsity may stall explanations (Strevens [2008]). Our proposal sidesteps such concerns by stripping out the problematic elements of falsity altogether. By limiting the verticalized theory's domain, we reinterpret it as straightforwardly true concerning a new, more abstract subject matter. This means we can avoid treating our old theories as merely useful fictions—which would jar with the central role they take both in practice and in explanation.

5. Theoretical Relicts

Given our conception of reduction as construction, reduction is a relationship between theories and is invariably a somewhat holistic business. Accordingly, we have thus far focused on reduction relations between theories, and we have taken theory-to-theory reductions as our primary target for verticalization. We have avoided making claims about reduction of individual entities to other individual entities.¹³ The model of reduction we have been working with requires us to recover the laws, variables, and equations—in other words, the nomological structure—of T_t from the underlying theory, or successor

¹³ We take it that the mereological view of levels and reduction given by Oppenheim and Putnam ([1958]) has been thoroughly repudiated by practice-informed accounts of reduction; see (Dupre [1993]; Potochnik [2017]).

theory, T_b . Achieving this reduction may require the help of various resources: boundary conditions, approximations and idealizations, and collective degrees of freedom.

Where does this leave the entities of reduced theories? Answering this question is complex, not least because there is no consensus about how to ‘read off’ the entities of a scientific theory. In other words, there are a variety of different forms of scientific realism; for a sample, see (Chakravartty [1998]; Psillos [1999]; Massimi [2004]; Ladyman and Ross [2008]). Here we aim to bracket questions about the relationships between theories and their characteristic entities so far as is possible, and instead focus on the consequences of reduction for the entities associated with our different scientific theories.

What we are calling vertical and horizontal reduction have typically been regarded as having different ontological consequences. In the case of horizontal reduction, the characteristic entities of superseded theories—theoretical relicts—are typically relegated to a status of ‘useful fictions’ at best. But in the case of vertical reduction, this is not the contemporary default: even if the theory of viruses can be completely understood in terms of amino acids, viruses are not eliminated from our scientific worldview. Vertical inter-theoretic reduction need not lead to elimination (Ney [2008]). When then should we adopt a realist attitude towards theoretical entities of higher-level theories? Some radical eliminativists do answer: never.¹⁴ But most realists adopt some more nuanced answer that is selective about candidate higher-level entities. Typically, the selection criteria employed have a link to explanation: broadly speaking, entities earn their keep through playing a role in our best explanations.

Our account has an unorthodox consequence: some theoretical relicts, entities of previously successful theories, are vindicated as higher-level entities. That is, they have the same status as more familiar higher-level entities such as gases, viruses, and phonons. After the domain restriction and reinterpretation, T_t is no longer an old, tainted theory (albeit one that is approximately true). Instead, within a limited domain, we can reinterpret the theory as true—but about some more abstract subject matter. Not only does this mean that any element of falsity is expunged—thus allowing the theory to offer genuine explanations¹⁵—but the old theory can now be considered a higher-level theory. Thus,

¹⁴ These include (Horgan and Potrč [2008]; French [2020]).

¹⁵ Recall that we are assuming an explanans must be true to qualify as an explanation; compare (Cartwright [1983], p. 91).

the entities of the verticalized theory—theoretical relicts—are now higher-level entities, and so part of the scientific realist's menagerie.

One clear example of a vindicated theoretical relict is space (as opposed to spacetime). Since the widespread adoption of the spacetime interpretation of special relativity, ordinary three-dimensional space has not been regarded as a fundamental ingredient of physical reality. Minkowski ([1952], p. 75) famously declared: 'From henceforth, space by itself, and time by itself, have vanished into the merest shadows and only a kind of blend of the two exists in its own right'. But spatial distances and temporal durations are not treated as mere shadows: they are ubiquitous in the higher-level sciences. This is easily accommodated by the verticalization procedure we have described: Newtonian mechanics, along with its relict entities space and time, is vindicated as correctly describing low-relative-velocity motion, rather than a mere fiction.

Here we can make a connection to the emergence literature.¹⁶ Spacetime is not considered to be fundamental in many theories of quantum gravity¹⁷; instead when general relativity is recovered (that is, horizontally reduced) from a theory of quantum gravity, it is immediately rehabilitated as an effective theory true within its domain.¹⁸ Spacetime is understood as emergent. In an analogous move, should we consider space and time to be emergent? We think so. In fact, the link to emergence allows us to put our answer to our motivating puzzle in slogan form: a good relict is an emergent relict. Here 'good' means 'vindicated': the relict entity is a characteristic entity of a theory whose reduction relation should be verticalized. In this way, our framework fills out a remark made in passing by Ladyman ([2018]): 'Entities that are now regarded as emergent are also often the entities of past theories'.

Recall from section 4 that our framework is suitably flexible that all horizontal reductions could in principle be verticalized—since to be accepted in the first place, they would need some element of empirical success, and the domain could just be limited to those instances, however few and miscellaneous they might be. But we added the condition that we should verticalize if and only if T_i is not strictly dominated by T_b with respect to

¹⁶ The type of emergence on the table here is weak emergence (Crane [2001]; Chalmers [2006]; Bedau and Humphreys [2008]; Wilson [2010]).

¹⁷ See, for instance, (Wüthrich et al. [2021]).

¹⁸ For further discussion, see (Crowther [2018]) and (Wallace [unpublished-b]) on general relativity as an effective theory.

explanatory power. This permits a particularly direct connection to a recent account of emergence due to Knox ([2016]). Knox and Franklin ([2018]) apply Knox's account to argue that phonons should be considered emergent since—within a particular domain and timescale—they are both robust and a source of novel explanatory value. Novel explanatory value is precisely what is required to vindicate a verticalized theory; it is typically associated with autonomous nomological structure in the higher-level sciences.¹⁹ In such cases, the nomological structure of T_t can be recovered from T_b , but nonetheless, in the domains where T_t is successful, the further fine-grained details of T_b do not matter.

The take-home message of this section is that when a reduced theory, T_t , still offers some genuine explanatory advantage, then it can be coherently reinterpreted as a higher-level theory—and then, we argue, a scientific realist should adopt the same attitudes to its entities, its theoretical relicts, as they do for other non-fundamental theories. Such theoretical relicts are not merely useful fictions but concrete scientific kinds.

6. How Not to Reify Theoretical Inconsistencies

Does the proliferation of entities at higher levels, and of true higher-level theories characterizing their behaviour, threaten pathology or inconsistency? Saatsi ([2022]), responding to a defence by Egg ([2021]) of a position that in some ways resembles our own, has recently argued that recovering the literal truth of superseded theories, even in a restricted domain, will lead to inevitable conflicts between theoretical claims of the newer and older theories. For example, Saatsi suggests that Newtonian gravitational theory and general relativity will inevitably end up making inconsistent claims about the concept 'gravitational force'. Saatsi worries, rightly in our view, that such inconsistency would play havoc with intertheoretic relations, multilevel models, and ultimately with the coherence of scientific realism. Our account, however, has the resources to avoid this catastrophe.

While the framework we endorse does keep open the possibility of reified inconsistencies—that is, our view doesn't somehow rule them out by definition—such inconsistencies are substantively precluded by the criteria we impose on verticalization. First, we

¹⁹ For discussion, see (Robertson [forthcoming]). This autonomy condition is also connected to Wilson's ([2010]) account of weak emergence, which involves eliminations in the number of degrees of freedom of the higher level compared to the lower level.

highlight again our stringent starting point: a successful horizontal reduction is a requirement of the first step in our verticalization procedure. Insofar as the new theory does indeed explain the success of the old theory, the new theory and the successful components of the old theory cannot disagree. This secures, post-verticalization, the supervenience of the higher-level theory on the lower-level theory. The domain-restriction steps of the verticalization procedure ensures that the components of the old theory that risk conflict with the new theory are factored out from the verticalized higher-level description; potential inconsistencies are eliminated by verticalization before they can be reified.

We hope that the role played in our account by the modal account of subject matters in implementing the domain-restriction step is now fully clear. In section 3 we argued that the modal account of subject matters provides a sense in which theories at different levels are about different phenomena and hence capable of having genuinely distinctive explanatory profiles. It also encodes the supervenience requirement between higher- and lower-level descriptions, thus forming a key part of our overall solution to Saatsi's problem of reified theoretical inconsistency.

The result is a levelled picture, where higher- and lower-level theories characterize different aspects of the phenomena at different levels. Our conditions of successful horizontal reduction, domain restriction and approximation provide assurance that the level-specific claims complement rather than clashing. On this approach, Galilean spacetime (for example) is no longer understood as merely a useful fiction; we understand Galilean spacetime as a description of spacetime structure at a 'classical' level of abstraction. Likewise, Newtonian gravity and general relativity characterize different structural features of gravitation located at different levels of abstraction.

We take it that as scientists (and causal reasoners more generally) we flexibly restrict our domains of quantification to include features from the right level to support our descriptive and explanatory goals. Thus 'gravitational force' is context-dependent, picking out different structure at different levels of description when used in different theoretical contexts. We don't need to acknowledge any legitimate context in which there is more than one gravitational force, but nor is there any unique level of description at which all gravitational structure resides. This strategy is philosophically conservative: it avoids any need to relativize truth or existence to levels, and instead draws on the well-understood semantic machinery of contextual domain restriction.

A properly verticalized reduction gives us a clear picture of how the higher-level theory depends on the underlying theory. Such a clear picture is lacking in the quantum-mechanical case on which Saatsi's criticism of Egg focuses,²⁰ but it is present in other examples that Saatsi discusses. These include some examples that play an important role in our argument, Newtonian mechanics and geometric optics; here our view does come up directly against Saatsi's objections to effective realism. Our overall response to Saatsi's argument in these cases is to hold the line on our core proposal, but to offer an alternative conceptualization of what Saatsi thinks will be missing from any account like ours.

Saatsi argues that a too-promiscuous effective scientific realism will be unable to make sense of the theoretical progress made precisely when we discover that certain entities (Newtonian forces, or light rays for example) are 'merely effective' and—so goes the line of thought—do not really exist.²¹ To some extent we bite the bullet here: within our framework we can offer alternative ways of making sense of what we discover when scientific progress relegates the status of the old theory's entities. For example, in the context of the reduction of geometric optics to wave optics, we regard light rays in geometric optics as non-fundamental emergent entities with a limited domain of applicability. This is in itself a significant metaphysical downgrade from their previous status, where they were assumed (a) to be at least relatively fundamental and (b) to be present in any scenario that featured any light whatsoever. In other words, we reject Saatsi's assumption that ontological elimination of theoretical entities constitutes an essential part of explanatory progress over instances of theory change—even radical theory change. Of course, our proposal does allow for cases of elimination of theoretical entities (in the case of theories not horizontally reduced, or where the horizontally reduced theory lacks any distinctive explanatory value)—but we see such cases as less central than Saatsi does.

In sum: The problem of reified theoretical inconsistencies is defanged by only considering successful reductions, and the metaphysical concern about entities with inconsistent properties is overcome by appealing to an underlying multi-levelled metaphysics

²⁰ Accordingly, we agree with Saatsi (*contra* Egg) that textbook 'orthodox' quantum mechanics is a poor candidate for an effective-realist rehabilitation.

²¹ This point should not be confused with so-called fictional forces in classical mechanics, such as 'centrifugal forces'. The latter are designated fictional not because they belong to a superseded theory but because they appear only in frames rotating with respect to an inertial frame.

and suitable contextual domain restriction. This approach generically avoids the possibility of conflicts between successfully verticalized theories and their underlying theories. Enabling this strategy for avoiding theoretical conflict, we take it, is one of the main virtues of 'levels talk'.

7. Conclusion

We began by distinguishing horizontal from vertical reduction. A horizontal relation shows how the old theory approximates the new theory, while a vertical reduction shows how the higher-level theory abstracts away from the lower-level theory. Approximations can be re-interpreted as abstractions, enabling horizontal reductions to be verticalized. The older theory is converted into a verticalized theory by being assigned a more constrained subject matter: the verticalized theory is restricted to apply in some narrower domain than that of the new theory, and it is reinterpreted as describing reality on a coarser-grained level. Verticalization turns old theories into effective theories and old theoretical entities into higher-level entities.

Verticalization is apt just when it gives rise to higher-level theories that are non-redundant in our best overall explanatory account of the world; that is, when they offer the best explanation for some high-level explananda. Verticalization rehabilitates the old theory: from being merely wrong, to being wrong as an account of the underlying details but right as an account of the emergent features.

When an old theory is vindicated through verticalization, so are the entities that play non-redundant explanatory roles in the resulting verticalized theories. Thus the theoretical entities of old theories—the 'theoretical relicts'—are revived as distinctively explanatorily powerful higher-level entities. This resolves our puzzle of theoretical relicts. Scientific realists should take these entities seriously only when they are emergent: when they figure in explanations provided by some aptly verticalized and distinctively explanatorily powerful higher-level theory.

Katie Robertson and Alastair Wilson

Acknowledgments

We are very grateful for invaluable comments on drafts of this paper to members of the FraMEPhys team—especially Nick Emmerson, Joaquim Giannotti, Mike Hicks, and Noelia Iranzo Ribera—to Alex Franklin, Miranda Rose, and Scott Sturgeon, and to two anonymous reviewers for this journal. We would also like to thank audiences at the Birmingham WiP seminar, the British Society for Philosophy of Science annual conference, the Bergen Philosophy of Science workshop, and the FraMEPhys/MetaScience joint workshop. This work forms part of the project A Framework for Metaphysical Explanation in Physics (FraMEPhys), which received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 757295). Funding for Alastair Wilson's contribution was also provided by the Australian Research Council (grant agreement no. DP180100105). Funding for Katie Robertson's contribution was also provided by the Leverhulme Trust Early Career Fellowship scheme.

Katie Robertson
Department of Philosophy
University of Birmingham
Birmingham, UK
k.e.robertson@bham.ac.uk

Alastair Wilson
Department of Philosophy
University of Birmingham
Birmingham, UK
Department of Philosophy
Monash University
Melbourne, Australia
a.j.wilson@bham.ac.uk

References

- Batterman, R. [1995]: 'Theories between Theories: Asymptotic Limiting Intertheoretic Relations', *Synthese*, **103**, pp. 171–201.
- Batterman, R. [2001]: *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*, Oxford: Oxford University Press.
- Bedau, M. and Humphreys, P. [2008]: *Emergence: Contemporary Readings in Philosophy and Science*, Cambridge, MA: MIT Press.
- Bickle, J. [1998]: *Psychoneural Reduction: The New Wave*, Cambridge, MA: MIT Press.
- Bird, A. [1998]: *Philosophy of Science*, Montreal and Kingston: McGill-Queen's University.
- Bird, A. [2000]: *Thomas Kuhn*, Chesham: Acumen.
- Bogen, J. and Woodward, J. [1988]: 'Saving the Phenomena', *The Philosophical Review*, **97**, pp. 303–52.
- Butterfield, J. [2011a]: 'Emergence, Reduction and Supervenience: A Varied Landscape', *Foundations of Physics*, **41**, pp. 920–59.
- Butterfield, J. [2011b]: 'Less Is Different: Emergence and Reduction Reconciled', *Foundations of Physics*, **41**, pp. 1065–135.
- Carroll, S. [2021]: 'Consciousness and the Laws of Physics', *Journal of Consciousness Studies*, **28**, pp. 16–31.
- Cartwright, N. [1983]: *How the Laws of Physics Lie*, Oxford: Oxford University Press.
- Chakravartty, A. [1998]: 'Semirealism', *Studies in History and Philosophy of Science Part A*, **29**, pp. 391–408.
- Chalmers, D. [2006]: 'Strong and Weak Emergence', in P. Clayton and P. Davies (eds), *The Re-emergence of Emergence*, Oxford: Oxford University Press, pp. 244–56.
- Churchland, P. M. [1981]: 'Eliminative Materialism and the Propositional Attitudes', *Journal of Philosophy*, **78**, pp. 67–90.
- Churchland, P. S. [1986]: *Neurophilosophy: Toward a Unified Science of the Mind/Brain*, Cambridge, MA: MIT Press.
- Crane, T. [2001]: 'The Significance of Emergence', in B. Loewer and C. Gillett (eds), *Physicalism and Its Discontents*, Cambridge: Cambridge University Press.
- Crowther, K. [2018]: 'What Is the Point of Reduction in Science?', *Erkenntnis*, **85**, pp. 1437–60.
- Dewar, N. [2019]: 'Supervenience, Reduction, and Translation', *Philosophy of Science*, **86**, pp. 942–54.
- Dupré, J. [1993]: *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Cambridge, MA: Harvard University Press.
- Esfeld, M. and Deckert, D. [2017]: *A Minimalist Ontology of the Natural World*, London: Routledge.
- Egg, M. [2021]: 'Quantum Ontology without Speculation', *European Journal for Philosophy of Science*, **11**, available at <doi.org/10.1007/s13194-020-00346-1>.
- Franklin, A. and Knox, E. [2018]: 'Emergence without Limits: The Case of Phonons', *Studies in History and Philosophy of Modern Physics*, **64**, pp. 68–78.

- Franklin-Hall, L. R. [2016]: 'High-Level Explanation and the Interventionist's "Variables Problem"', *British Journal for the Philosophy of Science*, **67**, pp. 553–77.
- French, S. [2014]: *The Structure of the World: Metaphysics and Representation*, Oxford: Oxford University Press.
- French, S. [2020]: *There Are No Such Things as Theories*, Oxford: Oxford University Press.
- Frigg, R. and Werndl, C. [2019]: 'Statistical Mechanics: A Tale of Two Theories', *The Monist*, **102**, pp. 424–38.
- Gardner, M. [1970]: 'The Fantastic Combinations of John Conway's New Solitaire Game "Life"', *Scientific American*, **223**, pp. 120–23.
- Hacking, I. [1982]: 'Experimentation and Scientific Realism', *Philosophical Topics*, **13**, pp. 71–87.
- Hartmann, S. [2002]: 'On Correspondence', *Studies in History and Philosophy of Modern Physics*, **33**, pp. 79–94.
- Hawke, P. [2018]: 'Theories of Aboutness', *Australasian Journal of Philosophy*, **96**, pp. 697–723.
- Horgan, T. and Potrč, M. [2008]: *Austere Realism: Contextual Semantics Meets Minimalist Ontology*, Cambridge, MA: MIT Press.
- Jones, R. [1991]: 'Realism about What?', *Philosophy of Science*, **58**, pp. 185–202.
- Koertge, N. [1969]: 'The General Correspondence Principle: A Study of Relations between Scientific Theories', Ph.D. Thesis, University of London.
- Kim, J. [1998]: *Mind in a Physical World: An Essay on the Mind–Body Problem and Mental Causation*, Cambridge, MA: MIT Press.
- Kitcher, P. [1981]: 'Explanatory Unification', *Philosophy of Science*, **48**, pp. 507–31.
- Kitcher, P. [1984]: '1953 and All That: A Tale of Two Sciences', *Philosophical Review*, **93**, pp. 335–73.
- Kitcher, P. [1993]: *The Advancement of Science*, Oxford: Oxford University Press
- Knox, E. [2016]: 'Abstraction and Its Limits: Finding Space for Novel Explanation', *Noûs*, **50**, pp. 41–60.
- Ladyman, J. [2011]: 'Structural Realism versus Standard Scientific Realism: The Case of Phlogiston and Dephlogisticated Air', *Synthese*, **180**, pp. 87–101.
- Ladyman, J. [2018]: 'Scientific Realism Again', *Spontaneous Generations*, **9**, pp. 99–107.
- Ladyman, J. and Ross, D. [2008]: *Every Thing Must Go: Metaphysics Naturalized*, Oxford: Oxford University Press.
- Lewis, D. [1980]: 'Mad Pain and Martian Pain', in N. Block (ed.), *Readings in the Philosophy of Psychology I*, Cambridge, MA: Harvard University Press, pp. 216–22.
- Lewis, D. [1988a]: 'Statements Partly about Observation', *Philosophical Papers*, **17**, pp. 1–31.
- Lewis, D. [1988b]: 'Relevant Implication', *Theoria*, **54**, pp. 161–74.
- List, C. [2014]: 'Free Will, Determinism, and the Possibility of Doing Otherwise', *Noûs*, **48**, pp. 156–78.

- List, C. [2019]: 'Levels: Descriptive, Explanatory, and Ontological', *Noûs*, **53**, pp. 852–83.
- List, C. and Menzies, P. [2009]: 'Nonreductive Physicalism and the Limits of the Exclusion Principle', *Journal of Philosophy*, **106**, pp. 475–502.
- List, C. and Pivato, M. [2015]: 'Emergent Chance', *The Philosophical Review*, **124**, pp. 119–15.
- Massimi, M. [2004]: 'Non-defensible Middle Ground for Experimental Realism: Why We Are Justified to Believe in Colored Quarks', *Philosophy of Science*, **71**, pp. 36–60.
- Minkowski, H. [1952]: 'Space and Time' in H. Lorentz, A. Einstein, H. Minkowski and H. Weyl (eds), *The Principle of Relativity: A Collection of Original Memoirs on the Special and General Theory of Relativity*, New York: Dover, pp. 75–91.
- Menzies, P. and List, C. [2010]: 'The Causal Autonomy of the Special Sciences', in C. McDonald and G. McDonald (eds), *Emergence in Mind*, Oxford: Oxford University Press, pp. 108–29.
- Myrvold, W. [2020]: "'— It would be possible to do a lengthy dialectical number on this;'" *Studies in History and Philosophy of Modern Physics*, **71**, pp. 209–19.
- Nagel, E. [1935]: 'The Logic of Reduction in the Sciences', *Erkenntnis*, **5**, pp. 46–52.
- Nagel, E. [1949]: 'The Meaning of Reduction in the Natural Sciences', in R. C. Stouffer (ed.), *Science and Civilization*, Madison: University of Wisconsin Press, pp. 99–135.
- Nagel, E. [1961]: *The Structure of Science: Problems in the Logic of Explanation*, New York: Harcourt, Brace and World.
- Nagel, E. [1970]: 'Issues in the Logic of Reductive Explanations', in H. E. Kiefer and K. M. Munitz (eds), *Mind, Science, and History*, Albany: SUNY Press, pp. 117–37.
- Ney, A. [2008]: 'Reductionism', *Internet Encyclopedia of Philosophy*, available at iep.utm.edu/red-ism/.
- Oppenheim, P. and Putnam, H. [1958]: 'Unity of Science as a Working Hypothesis', *Minnesota Studies in the Philosophy of Science*, **2**, pp. 3–36.
- Potochnik, A. [2017]: *Idealization and the Aims of Science*, Chicago, IL: University of Chicago Press.
- Psillos, S. [1999]: *Scientific Realism: How Science Tracks Truth*, London: Routledge.
- Putnam, H. [1975]: 'Philosophy and Our Mental Life', in his *Philosophical Papers 2: Mind, Language, and Reality*, Cambridge: Cambridge University Press, pp. 291–303.
- Railton, P. [1978]: 'A Deductive-Nomological Model of Probabilistic Explanation', *Philosophy of Science*, **45**, pp. 206–26.
- Railton, P. [1981]: 'Probability, Explanation, and Information', *Synthese*, **48**, pp. 233–56.
- Robertson, K. [2019]: 'Reductive Aspects of Thermal Physics', Ph.D. Thesis, University of Cambridge.
- Robertson, K. [2020]: 'Asymmetry, Abstraction, and Autonomy: Justifying Coarse-Graining in Statistical Mechanics', *British Journal for the Philosophy of Science*, **71**, pp. 547–79.
- Robertson, K. [forthcoming]: 'Autonomy Generalised; or, Why Doesn't Physics Matter More?', *Ergo*.

- Rosaler, J. [2015]: 'Local Reduction in Physics', *Studies in History and Philosophy of Modern Physics*, **50**, pp. 54–69.
- Rosaler, J. [2017]: 'Reduction as an *a posteriori* Relation', *British Journal for the Philosophy of Science*, **70**, pp. 269–99.
- Saatsi, J. [2022]: '(In)effective Realism?', *European Journal for Philosophy of Science*, **12**, available at <doi.org/10.1007/s13194-021-00441-x>.
- Schaffner, K. [1967]: 'Approaches to Reduction', *Philosophy of Science*, **34**, pp. 137–47.
- Strevens, M. [2008]: *Depth: An Account of Scientific Explanation*, Cambridge, MA: Harvard University Press.
- Suppes, P. [1967]: 'What Is a Scientific Theory?', in S. Morgenbesser (ed.), *Philosophy of Science Today*, New York: Basic Books, pp. 55–67.
- van Gulick, R., [2001]: 'Reduction, Emergence, and Other Recent Options on the Mind/Body Problem: A Philosophic Overview', *Journal of Consciousness Studies*, **8**, pp. 1–34.
- van Riel, R. and van Gulick, R., [2019]: 'Scientific Reduction', in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, available at <plato.stanford.edu/archives/spr2019/entries/scientific-reduction/>.
- Wallace, D. [2012]: 'The Arrow of Time in Physics', in A. Bardon and H. Dyke (eds), *A Companion to the Philosophy of Time*, Chichester: Wiley-Blackwell.
- Wallace, D. [unpublished-a]: 'The Logic of the Past Hypothesis', available at <philsci-archive.pitt.edu/8894/>.
- Wallace, D. [unpublished-b]: 'Quantum Gravity at Low Energies', available at <philsci-archive.pitt.edu/20049/1/low_energy_qg.pdf>.
- Wallace, D. [unpublished-c]: 'Stating Structural Realism: Mathematics-First Approaches to Physics and Metaphysics', available at <philsci-archive.pitt.edu/20048/1/semantic.pdf>.
- Weslake, B. [2010]: 'Explanatory Depth', *Philosophy of Science*, **77**, pp. 273–94.
- Williams, P. [2019]: 'Scientific Realism Made Effective', *British Journal for the Philosophy of Science*, **70**, pp. 209–37.
- Wilson, J. [2010]: 'Non-reductive Physicalism and Degrees of Freedom', *British Journal for the Philosophy of Science*, **61**, pp. 279–311.
- Woodward, J. [2008]: 'Mental Causation and Neural Mechanisms', in J. Hohwy and J. Kallestrup (eds), *Being Reduced: New Essays on Reduction, Explanation, And Causation*, Oxford: Oxford University Press, pp. 218–62.
- Woodward, J. [2021]: 'Explanatory Autonomy: The Role of Proportionality, Stability, and Conditional Irrelevance', *Synthese*, **198**, pp. 237–65.
- Wüthrich, C., Le Bihan, B., and Huggett, N. [2021]: *Philosophy beyond Spacetime*, Oxford: Oxford University Press
- Yablo, S. [1992]: 'Mental Causation', *The Philosophical Review*, **101**, pp. 245–80.
- Zeh, H. D. [2007]: *The Physical Basis of the Direction of Time*, Berlin: Springer.