Against Rigour: In Favour of Informal Standards for Reduction in Philosophy of Physics

Alexander Franklin*

March 5, 2026

Intended for *The Philosophy of Rigour*, published by Routledge, edited by Dean Rickles and Karim Thébault.

Abstract

Many would regard rigour as self-evidently a virtue in philosophical theorising. By considering two vignettes from recent debates over reduction in philosophy of physics I'll argue that the pursuit of rigour has distracted philosophers from the questions that motivated their research: in both contexts I'll suggest that less rigorous analysis would better enable such questions to be satisfactorily answered. I'll claim that these case studies exemplify aspects of the broadside against rigour developed in Wilson (2021).

> [S]tandards of surface "rigor" based upon a faulty conception of "science" hinder productive investigative effort within many corners of the philosophical world, commonly inducing an overprizing of ersatz mathematization
>
> Wilson (2021, p. 5)

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

1 Introduction

What is the point of rigour? What's it for? And whose interests does it serve? Certainly, there's some all-else-being-equal added value to rigorous argument and analysis in philosophy. The goal of this chapter is to discuss some of the trade-offs involved in upping the rigour level in certain philosophical discussions. My examples come from debates over inter-theoretic relations but my inspiration for much of this analysis comes from the work of Mark Wilson who considers a number of other examples in philosophy, and so I have some grounds for thinking the argument generalises.

I make two claims in this paper: first that the pursuit of rigorous analysis can distract from the philosophical projects that are the explicit ends of some of those philosophers discussed – that is, rigour serves as a poor regulative ideal; second that there's a trade-off between rigour and generality but that this is a subtle trade-off: it's my view that much interesting philosophical work takes place in a sweet spot between the very abstract universal and the absolutely particular. Rigorous analysis is often best suited to one extreme or the other. While the mathematical physicist and logically inclined philosopher can make absolutely general claims with formal reasoning, and can also reason and argue about the details of an absolutely specific domain, such reasoning can be ill-suited to the spotting of relations and patterns between different domains, the noticing and adducing of common systems of reasoning in different contexts, and mid-range abstractions. This paper is an exercise in demonstrating that that mid-range context is fruitful but I'll claim that rigorisation of the analysis of this context can mislead.

As the subtitle makes clear, my focus is on instances of rigour in discussions of reduction in the philosophy of science and philosophy of physics. Rigour and exhortations to greater rigour take two forms in the contexts I'm discussing here: in part they involve the explicit specification of criteria that reduction ought to satisfy – see the Nagel and Nickles formal models of reduction; in philosophy of physics they also involve attempts to convert the kinds of derivation used in physics textbooks and by physicists into the more mathematised derivations that remove limits or attempt to satisfy certain elevated mathematical standards.

My arguments for these claims come in two parts. First they involve the examination of the two models of reduction. In each case I suggest ways in which the demonstration of reduction, insofar as it's rigorous, either falls short in its generality or leads to a debate that's not fruitful. I suggest that a less formal approach focussed on understanding and explanation would better serve the goals and interests of the philosophers engaged in the discussion. Second I examine explanations as to why our concepts and notions may not be susceptible to increasing rigorisation: increasing the precision of our concepts may be at odds with their capacity accurately to model the world.

This latter point relates closely to Wilson's (2021) claims that the standards of rigour employed in the sciences are imported to philosophy in a way that cannot serve the same ends. In part this is because philosophy is often engaged in much more general or universal questions than those discussed in any particular science. As such, even if rigour in a specific context is salutary, when one attempts to generalise one takes a concept with a specific meaning in a domain and pulls it away from the context where it gained that meaning, this can often lead to inaccuracy or false claims. Insofar as one seeks to retain rigour one generalises to the point where any applicability is lost (or at least hard to see).

In addition philosophers often mischaracterise scientists' interest in rigour and imagine that it's more important to the scientific enterprise than is observed by an analysis of the relevant work. This is a point found in Wilson (2021) but also emphasised in David Wallace's (2022). Both philosophers suggest that by paying attention to the locution of science and the derivational practices of scientists rather than formal and rigorous logical analysis one may make more philosophical progress. The thought is that philosophers aspire to greater rigour than scientists do and that going further is rarely achievable and often misleading where it appears to be available; I suggest that this is especially true in the study of inter-theoretic relations where we require our reasoning and modelling practices to traverse scientific domains.

Wilson's work focusses on the details of physicists' and philosophers' efforts in the 19th and early 20th centuries to construct rigorous foundations for physics, and he notes that these efforts were often frustrated. He suggests that the best way to understand classical physics (though we might well expect this to extend to quantum physics) is as a set of interconnecting submodels to which we have to pay special attention when moving between subdomains. The particular styles of rigorisation employed in various philosophical exercises that Wilson calls 'theory T thinking' have a tendency to universalise in a way that loses accuracy – for example one might be able to make a few very general claims about the quantum-classical limit, but the details will depend on each instance of that limit, as covered by decoherence theory, see §3. Wallace's work does something closely related: he notes the tension between philosophers' (especially metaphysicians') appeal to certain formal methods in contrast to the more slapdash methodology of practising physicists. Wilson and Wallace both note that the toolbox used by those scientists whose work achieves empirical success is rather different from the logical and mathematical methods employed by those seeking greater rigour, and it's not at all clear that the latter toolbox is adequate to describing the world – thus the tension between rigour and scientific applicability. One might of course note these issues and still aspire to rigorise such sloppy scientific thinking – that would be to aspire to greater rigour. In the following I'll take instances where that's occurred and suggest that doing so isn't as fruitful as some might have expected. Themes from both Wilson and Wallace therefore explain why greater rigour should not be the aim of many philosophical enterprises.

The overall aim of this paper is to offer some resistance to what seems to be a trend in philosophy of physics and philosophy of science, a trend which has some sociological reasons in its favour, see Jingyi Wu's unpublished. Despite the provocative title I'm not against rigour in general or in all its forms; rigour certainly has a function in discourse, but more rigour isn't always desirable and I hope to demonstrate that rigour can lead us astray in the pursuit of truth, understanding, and well justified claims about the nature of science and the world it purports to describe. By analogy: Railton (1981) famously argued that the best explanations involve the most possible detail, and responses due to Batterman (1992), Craver and Kaplan (2018), and Yablo (1992), (building on Putnam (1975)) among many others pointed out that more details aren't always better, and in some contexts extra details detract from explanatory generality and goodness – in fact sometimes more details lead to the acceptance of certain false counterfactuals (e.g. that the pigeon wouldn't have pecked had the speck been scarlet rather than maroon).

Rigour is desirable because rigorous arguments are those where it's perfectly clear which claims one must give up in order to reject the conclusion. More generally, rigorous models offer apparently clearly expressed criteria by which one might successfully make a case for a given claim. Rigour, thus, is a byword for clarity, and clarity is of course desirable. Improvements in rigour, however, are not costless. The requirement of rigour might cost significant time or even make certain projects unachievable. And, as noted above, rigour may also limit generality or applicability. Lastly, rigour may lead to value capture, where a given set of aims is replaced by a subtly different, rigorously achievable set of aims, which no longer serve the philosophical goals that they replaced (cf. Nguyen (2020)): debates may be captured by back and forths over the satisfaction of the criteria of the models of reduction rather than the heuristically expressed philosophical goals. It's for these reasons that rigour is often not desirable.

There's a tension in the background. It's commonplace for many to take the unsusceptibility to rigorisation of much physics practice as grounds for denying realism. This may be thought of a physics-in-practice worry about being able to read off claims about the world from our models; see e.g. Potochnik (2017) and Bokulich (2008), and many others. Wilson and Wallace argue that we should not deny that our models and theories are about the world, rather that there might be a limit to what we can say beyond the claim that our models represent the world accurately in given domains. It's not that there are no contexts in which extra rigour is desirable or achievable but that insofar as rigorisation fails we ought not to conclude that our models are purely a scientific artefact, it might just be that the best way to understand the world isn't through the philosophers' logic or the universal, abstracted, and rigorous tools of the mathematical physicist but through the messy and domain restricted language of theoretical physics. Note that the work of the physics in practice philosophers is still valuable – as noted below such work often brings out idealisations, falsehoods, and invalid assumptions that require detailed philosophical analysis to justify, it's just that the conclusion that a given model doesn't represent the world on account of such idealisations may be denied.

The case studies I'm exploring in this paper engage with questions of inter-theoretic relations. Reductionism is the thesis that facts about fundamental physics are in some sense sufficient to determine all other facts. Of course, this thesis takes many forms and is much contested. But in my view debates over reduction and reductionism are enlightening and philosophically fruitful notwithstanding this lack of precise definition. How such debates proceed, at least in many cases, is that some exemplar that's especially difficult to reduce is posited in the form of an explanatory gap – that is a phenomenon that robustly appears in the world at some level and yet seems not to be explainable from or perhaps not even consistent with the more fundamental science. The business of seeking to account for that phenomenon, to explain it, or to demonstrate its consistency with the more fundamental description is the business of reduction. The truth or falsity of reductionism is established on an evaluation of the totality of such case studies. Reductionism implies that all the explanatory gaps can be closed, and it fails if some are resistant to being filled in even in principle . My contention in this paper is that such debates would fare better (i.e. would deliver more insights, shed more light, and be more swiftly resolved) were the more rigorous and formal requirements on reduction to be dropped in favour of the much looser goals of closing explanatory gaps and seeking to understand the features of the less fundamental description in more fundamental terms.

The following discussion engages with two cases in which I suggest that overly formal and rigorous accounts of what's required by reduction are prone to value capture, where satisfying the criteria of the account diverges from what's of primary philosophical interest and thus diverts the debate unhelpfully. Should we nonetheless try to say more precisely what reduction is and what it involves? I suspect that such sharpening may well fail systematically, and I'm not persuaded that more is required than that expressed in the paragraph above.

In each section that follows, I'll set out the model of reduction – respectively, Nagel and Nickles reduction – and discuss a vignette from the recent literature where I'll argue that the disvalue of philosophical rigour is evident. Not only do we never achieve perfect rigour, but we also lack a good basis to suppose that perfect rigour is achievable in principle.

2 Case I: Nagel Reduction and Phase Transitions

The Nagel model of reduction is rigorous in the sense that it purports to offer a logically valid arguments for the truth of reduction in each context to which it's applied. In this section I'll suggest reasons to think that the debate over reduction should not depend on the details to which the model itself draws focus. My primary claim is that, *pace* Dizadji-Bahmani, Frigg, and Hartmann (2010), the Nagel model should not serve as a regulative ideal for reduction.

The Nagel model starts out with an explanatory condition: that the lower-level theory explains aspects of the higher-level theory.

> Reduction, in the sense in which the word is here employed, is the explanation of a theory or a set of experimental laws established in one area of inquiry, by a theory usually though not invariably formulated for some other domain.
> [Nagel (1961, p. 338)]

This non-rigorous account of what reduction involves accords with the view defended in this paper. Unfortunately Nagel goes on to develop his account in line with the Deductive-Nomological theory of explanation, and thus expresses formal criteria that any reduction ought to satisfy. This account of explanation holds that a fact is explained if and only if that fact is deduced from the laws and further particular facts.

The Nagel model is often expressed as requiring two conditions, labelled 'connectability' and 'derivability': first, that there are relations, known as 'bridge laws', between terms used in higher-level and lower-level theories – paradigmatically, the temperature of a gas is related to the mean kinetic energy of the constituent particles of the gas; second, the laws of the higher-level theory must be derivable from the laws of the lower-level theory, where such derivations generally appeal to the connecting relations of the first condition.

Schaffner's contribution was to allow that the higher-level and lower-level theories may not exactly match in their predictions. Thus a reduction may still occur where one has merely derived an approximation to the relevant higher-level theory. As the following quotation summarises:

> T_F [lower-level theory] reduces T_P [higher-level theory] iff there is a corrected version T_P^* of T_P such that, (a) T_P^* is derivable from T_F given that the terms of T_P^* are associated via bridge laws with terms of T_F, and that (b) the relation between T_P^* and T_P is one of, at least, *strong analogy* (sometimes also 'approximate equality', 'close agreement', or 'good approximation').
> [Dizadji-Bahmani, Frigg, and Hartmann (2010, p. 398)]

Moreover, Schaffner (2013, p. 53) defines partial reductions in contexts where complete reductions are unavailable:

> partial reductions occur in the context of almost-complete reductions if and only if one or both of the generalized conditions of connectability (condition one) and derivability (condition two) of the GRR [GNS] model partially fail, or if the corrections and close-analogy conditions (conditions three and four) partially fail.

The distance between the original quote from Nagel and its explication can be attributed in part to Nagel's proximity to Logical Empiricism. Much of Wilson (2021) is engaged with an historical analysis that identifies many of the issues in contemporary metaphysics with a hangover from the formalisation and rigorisation projects of the early 20th century. Wilson claims that these projects were well motivated at the time and stemmed from the grand works of 19th century philosopher-physicists. The thought was that one could come up with a coherent foundation that, for example, resolved worries about the mysterious nature of the concept force in classical physics. Wilson describes how such projects failed for principled reasons – that in fact the different subfields of classical physics use 'force' in multiple distinct ways. The logical empiricists took up this project notwithstanding its issues and attempted to rigorise science more generally; as is well known that project failed, however parts of philosophy retain its vestiges and this is the subject of Wilson's critique: he argues that this is based on an attempt of philosophy to imitate science while philosophers have failed to notice that science gave up on such rigorisation long ago.

It's no coincidence that Nagel's rigorisation of reduction also stems from the logical empiricist tradition – I'd argue that it's misconceived in a similar way. As I'll discuss further below the very idea of a bridge law presupposes a systematic commonality between scientific terms in different contexts that often does not hold true. This is further assumed to take the form of a type-type identity between terms at different levels or scales.

I turn now to the case of phase transitions where I think that focus on Nagelian reduction and its bridge laws has distracted from the crux of the issues. It'd be much better to consider explanation of the phenomena rather than satisfaction of the formal criteria.

Phase transitions are characterised mathematically by a singularity and physically by a discontinuity: there's an abrupt change in physical behaviour between ferromagnets and paramagnets, or between liquids and gases. Starting with the phenomenology, gases that at macroscopic scales are homogeneous condense to form liquids that are similarly macroscopically homogeneous, and the transition between those two states seems to defy a straightforward understanding. The suggestion is that there's an explanatory gap between the macroscopic characterisations of the liquid and the gas and their molecular/atomic more fundamental description.

The puzzle seems first to have been raised in the philosophy literature in Liu (1999), and crystallised by Callender (2001). Since then, the anti-reductionist baton has been taken up most famously by Batterman (2005; 2010; 2011; 2013). Given the availability of derivations of the principles and relations of thermodynamics from those of statistical mechanics, anti-reductionism is advocated by putting pressure on the assumptions made in such derivations.

This putative explanatory gap can be sharpened by considering the mathematised description from physics. It's straightforwardly demonstrable that finite statistical mechanics cannot recover the singularities posited in the thermodynamical description of the phase transition.

The free energy can change only if work is done on or by a system. It is a function of the temperature (T), entropy (S), pressure (p), volume (V) and chemical potential (μ). An infinitesimal change in the (Helmholtz) free energy of an N particle system is written:

$$dF = -pdV - SdT + \sum_{i=1}^{N} \mu_i dN_i \tag{1}$$

First order phase transitions are defined in thermodynamics as corresponding to a non-analyticity (discontinuity) in the free energy. Thus the function F which is defined by integrating equation (1) will be smooth except at a phase transition, at which point equation (1) would cease to be uniquely defined, hence F would not be a smooth function.

The free energy is defined in statistical mechanics to be $\mathcal{F} = -k_B T \ln Z$. If we are to carry across the thermodynamic definition of a phase transition then it will be required that $\mathcal{F}$ be non-analytic – that it has a discontinuity in its first derivative. But the first derivative of $\mathcal{F}$ will be continuous: it is just the derivative of the logarithm of a finite sum of continuous functions. However, if we move to the thermodynamic limit, in which the system has infinitely many components, but a finite density, the first derivative may be discontinuous, and $\mathcal{F}_\infty$ may be non-analytic; see Styer (2004) for more details. Set up in this way, the thermodynamic definition of phase transitions does not have an analogue in finite statistical mechanics: $\mathcal{F}_N$ for $N \neq \infty$ will be analytic.

Batterman among others claim that insofar as reduction is viewed in Nagelian terms, this is a failure of reduction because the putative bridge law between the thermodynamic and statistical mechanical free energies requires the unphysical idealisation of an infinite N. Since bridge laws are supposed to tell us which quantities are to be identified this looks like a straightforward failure for reduction. Intriguingly Batterman (2010b, p. 176) does accept that "the divergences and singularities at critical phases are not genuine obstacles to some kind of general limiting (reductive?) relation between the theories after all" – I'll return to discuss limit-based reductions in §3.

Yet another approach to the puzzle of phase transitions employs the framework of symmetry breaking (both conceptually and in the mathematical representation): a condensing gas must break translational and rotational symmetries in order to form a liquid, as must a liquid to form a solid; likewise the (anti-)ferromagnetic phase breaks symmetries relative to the paramagnetic phase. This view of the puzzle of phase transitions is developed in Morrison (2006) where it's claimed that this is an instance of the failure of reduction because of the requirement of a top-down imposition of broken symmetry.

Does the more fundamental theory have the resources to account for the phase change that breaks the symmetries, or to make sense of the appearance of discontinuities in the higher-level theory? These are the questions and explanatory gaps that motivate anti-reductionist claims.

The answers to these questions are hard to answer, below I'll discuss some attempts. But first I think that it's instructive for the purposes of this paper to consider how this debate has played out in parts of the philosophy of physics literature that have focussed on Nagelian reduction.

My targets are Butterfield (2011), Butterfield and Bouatta (2012), Dizadji-Bahmani (2021), and Dizadji-Bahmani, Frigg, and Hartmann (2010). My principal claim is that the apparent explanatory gaps, expressed in one or other of the ways above is a genuine philosophical conundrum. The project of understanding how the world fits together from the bottom up (there are of course many ways to characterise the reductionist project) is worthwhile in part because it might evidence some general philosophical thesis about wholesale reductionism, strong emergence, or some weaker intermediate thesis, but also more narrowly because it relates to the specific project of understanding the world in all its local peculiarities. It's a puzzle that's worth engaging with if the world has certain properties in specific domains that cannot be immediately explained in some more fundamental terms. And the case of phase transitions seems to instantiate such a puzzle. My overall worry is that the Nagelian solutions are procedural fixes: they resolve the surface-level puzzle without addressing the deeper issues. Why this matters is that once one has stipulated bridge laws of a certain kind, the need for further philosophical enquiry will be obscured, and the potential for further valuable philosophical insights will be foreshortened.

So, while much of the literature does contribute towards a resolution of the issues raised by Batterman and others, and in fact, as detailed below, explanations are provided that may close the relevant explanatory gap, phrasing these in terms of Nagelian reduction is counterproductive. What that does is move the focus to questions over the nature of bridge laws and whether or not these contain illicit components. Instead, seeking to close explanatory gaps and provide how-possibly explanations does far more to address anti-reductionist claims head on, and insofar as such explanations are unavailable it's made clearer why reduction fails. Butterfield's work is exemplary of the issues raised here: he both provides exactly the kinds of explanations that demonstrate how sharp, discontinuous changes may happen in systems described with finite degrees of freedom (under his label 'Before') – thereby closing some explanatory gaps – but he also states that these should be thought of as Nagelian reductions and that this is established by a limiting bridge law; it's the latter that distracts from the crucial work he does with the former analysis.

As noted, Batterman is motivated by the distinct observable change at the phase transition. One might think that such macroscopic changes correspond to physical discontinuities, and may only be modelled by discontinuous mathematics. However, there is little reason to suppose that what we observe can only adequately be modelled by discontinuous mathematics as opposed to models which approximate the discontinuity. As Bangu (2009, 2015) argues, no measurements, however fine the precision, can distinguish between a true discontinuity and an arbitrarily steep gradient. As such this argument does not tell against the reductionist who wishes to interpret the infinite model as, merely, approximating the finite physics.

In this vein Butterfield (2011, p. 1128) notes: "there are physical quantities for finite models whose gradients grow without bound as $N \to \infty$". Butterfield gives the example of a magnetic phase transition in a ferromagnet: as the magnetic field rotates ($-H$ through 0 to $+H$), there will be a transition from spins aligned one way to spins aligned in the opposite orientation ($-M$ through 0 to $+M$). While this is modelled as discontinuous (graphically a step function), requiring an infinite idealisation, Butterfield points out that the mathematical description of the physical system will approach the discontinuity as the number of spins present is increased (the function will get steeper, thus looking more like a step function as $N \to \infty$). This is illustrated in figure 1.

> a larger N acts as a brake on the ferromagnet's response to the applied field increasing from negative to positive values (along the given axis). That is: the increased number of nearest neighbours means that the ferromagnet "lingers longer", has "more inertia", before the rising value of the applied field succeeds in flipping the magnetization from -1 to +1. More precisely: as N increases, most of the change in the magnetization occurs more and more steeply, i.e. occurs in a smaller and smaller interval around the applied field being zero. ...
>
> And this general picture of the approach to the $N \to \infty$ limits applies much more widely. In particular, very similar remarks apply to liquid-gas phase transition, i.e. boiling. There the quantity which becomes infinite in the $N \to \infty$ limit, i.e. the analogue of the magnetic susceptibility, is the compressibility, defined as the derivative of the density with respect to the pressure
>
> [Butterfield (2011, p. 1128)]

Insofar as Butterfield's claims may be derived from the bottom up, his arguments provide a reductive explanation for the appeal to the infinite limit: while it is acknowledged that the infinite model is more tractable, we would have a bottom-up explanation for how the finite system's properties are approximated by the infinite model. Moreover, Butterfield's claims can be clearly demonstrated on the Ising model which provides a coarse-grained model of a ferromagnet – thus there is good reason to expect that such reductive explanations can be provided for some systems which undergo first order phase transitions. These considerations and this kind of explanation form part of the project that is required to make progress

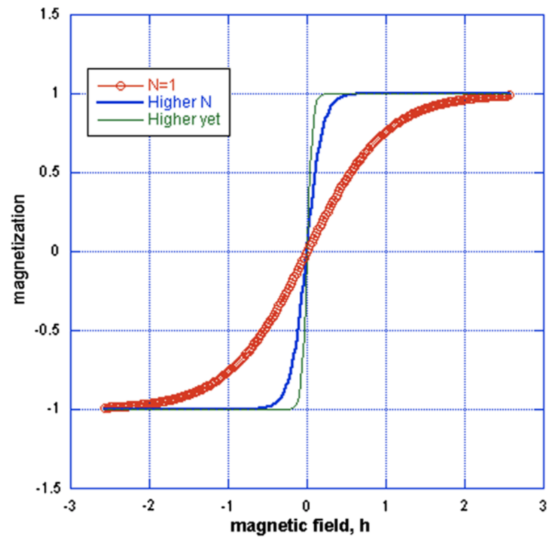

Figure 1: Kadanoff (2009, p. 783) demonstrates the effect of increasing N on magnetisation curves.

towards a detailed explanation of phase transitions from the bottom up. More work is needed to demonstrate that this applies to the liquid-gas case. Importantly, and in line with our theme, such work requires case by case analysis: Butterfield's quote above acknowledges this implicitly where he justifies the applicability to liquid-gas phase transitions in the second paragraph.

To flesh out his claim there: isothermal compressibility (κ) is defined as $\kappa_T = \frac{-1}{V}\left(\frac{\partial V}{\partial p}\right)_T$. This corresponds to how much the volume will change (∂V) with a given pressure change (∂p) at fixed temperature (T). The compressibility diverges for systems undergoing first order phase transitions because both phases are present and gases have far higher compressibility than liquids.

Where the compressibility diverges the system's volume may change by an arbitrarily large factor due to small changes in pressure. This might explain why the physics allows for abstraction away from particle number (N). Thus we have a story that provides a partial explanation: the processes which relate the available kinetic energy to the divergent compressibility explains the irrelevance of the particle number at a first order phase transition. Under what conditions does this number remain irrelevant? The compressibility only diverges while there is a mixture of liquid and gas, as such, the abstraction is only justifiable in this way while the system is undergoing a first order phase transition. Crucially, this will only work for sufficiently large systems where the $1/N \sim 0$ assumption is applicable.

This story gives us the flavour of how explanatory gaps of this kind can be resolved. If finite properties are well-approximated by infinite properties, as Butterfield shows, and the abstractions to move to the infinite model are justified, then little room is left for anti-reductionism. The question 'why the thermodynamic limit?' is addressed by observing that the thermodynamic limit allows for abstraction from lower-level details. The observation that abstraction is the core advantage of the thermodynamic limit is found in Knox (2016) and apparent from Batterman's focus on minimal model explanations, see Batterman and Rice (2014). Without accounting for the abstraction, we haven't fully resolved the puzzle of phase transitions. A reductionist should be in the game of justifying such lossy abstractions from the bottom up, and thus demonstrating that phase transitions are compatible with reduction. The empirical success of appeals to infinite N statistical mechanics illustrates the capacity for abstraction in our description of the world; while such a capacity remains unjustified, we cannot claim that the use of the thermodynamic limit has been fully explained from the bottom up.

What concerns me is that parsing this reductionist work in terms of the Nagel model requires a move to a level of generality that obscures the crucial physical explanations. In particular, the appeal to bridge laws is at a level of abstraction that precludes the context-relative closing of explanatory gaps.

However, Butterfield argues that we should understand his account as a Nagelian reduction, where the bridge law – $\mathcal{F}_N \iff \lim_{N\to\infty} \mathcal{F}_N$ – which connects thermodynamics to finite statistical mechanics, and which allows the derivation of the thermodynamical laws, involves an infinite limit; this builds on earlier work by Mainwood (2006). On the other hand, one might doubt whether such limiting bridge laws are sufficient for a reduction; an anti-reductionist who thinks that the singularity is more than calculationally essential is unlikely to be persuaded by this putative reduction for it involves the infinite limit – the very unphysical idealisation that's the subject of their worries. I am sceptical as to whether this question can be resolved by thinking about which bridge laws are permissible in the abstract; instead, given a limiting bridge law, one ought to ask whether that is justifiable from the bottom up and engage in the type of reasoning above.

What have we gained from all the ink spilt on this debate? I think interesting work has been developed on the various approximations and idealisations employed in scientific modelling. Batterman's work on the importance of robustness and autonomy – both of which essentially involve abstractions from lower-level details – is also insightful. But the purported central focus, the question of whether or not phase transitions are instances of irreducible physical phenomena is not conclusively settled until the (approximate) derivation is available and justifications of all the required abstractions have been set out.

My contention in this paper is that the discussion of the Nagel model is a distraction from the interesting physical details of the debate. Butterfield's emphasis on the behaviour of finite N systems before the limit is illuminating, but the attempt to integrate this with any bridge law is unhelpful: it could not convince any sceptic who wished to understand why the limit is needed.

Surely it would just be better to focus on explaining and understanding how it is that we end up with the higher-level description. If that cannot be explained then we have putative and empirically defeasible grounds for rejecting reduction in this case. But why would such discussions be aided by attempts at rigour in the formulation of the criteria for reduction?

I'll conclude discussion of this case study by developing two further reasons that we should be wary of Nagelian approaches to reduction: first I'll detail reasons for thinking that they mischaracterise the worldly metaphysics and thus there are principled reasons why bridge laws should be unavailable, and second I'll raise a response to the puzzle of phase transitions that the focus on Nagelian reduction may have played a role in obscuring. Note that further reasons to be sceptical of Nagelian approaches have been articulated by Batterman (2000) focussing on multiple realisation and Batterman (2013) with an emphasis on the homogenisation assumptions made by such models; these are respectively taken up from a reductionist (but anti-Nagelian) viewpoint in Franklin (2021) and Franklin (2024).

It's no accident that the advocates of Nagelian reduction focus on certain simple case studies – one finds reference to the temperature-mean kinetic energy relation, or mass in Newton vs. mass in special relativity to be commonplace, and even the more complex examples in Butterfield

(2011) are relatively idealised and simplified. What's sometimes unclear is whether such examples are to be taken as paradigms in the sense that we should expect the terms in science more generally to conform to this pattern of relation via bridge law, or whether they are acknowledged to be special cases. However Dizadji-Bahmani, Frigg, and Hartmann (2010, p. 410) are explicit: "[w]e have argued that GNS [Generalised Nagel Schaffner reduction] is alive and well, and that scientists involved in a reductionist research programme do the right thing if they take GNS as a regulative ideal."

It's my suggestion in this paper that this is a poor regulative ideal – that's in part because bridge laws are at odds with the true metaphysics of our physical systems. According to Dizadji-Bahmani et al. bridge laws are factual claims, not conventions corresponding to regularities, Humean or otherwise. "The first kind [of bridge law] associates basic entities of T_P and T_F with each other; they identify, for instance, light and electromagnetic radiation, electric currents and the flow of electrons, and gases and swarms of atoms" (Dizadji-Bahmani, Frigg, and Hartmann (2010, p. 404)). But is there really a regularity associating types of entities in this case? What, precisely, is a swarm of atoms? Is it just a collection of atoms arranged gas-wise? And, if so, what does the bridge law do for us?

If, rather, there is no type at the atomic level that corresponds to a gas one might think of the gas-swarm relation as a token-token relation, but the token-token links are remarkably fleeting. Given the robustness of gas states, there will be many distinct swarms that constitute the gas from one moment to the next. So the identification at any one time doesn't serve any scientific or philosophical ends. And once you move into the quantum context the few-many linkage is again much harder to make out, because the division into parts is lossy with respect to the entanglement relations; this is especially clear when one considers distributed processes such as decoherence – see §3.

Given that entity relations are not at all straightforward, property associations won't fare better. Again, we have paradigm cases that some think will generalise but in fact are special cases. For example, the temperature-mean kinetic energy doesn't hold for solids, see Robertson (2019).

Wallace (2022, pp. 15–16) develops worries of this ilk by considering how one could make formal sense of the identity between a swarm and a gas:

> One common answer . . . is to enrich our ontology with composites or mereological sums of particles . . . And in some cases mereology does not even superficially seem to fit the case. The quantum theory of vibrations in crystalline solids is normally described in terms of phonons – quantized particles of vibration. (For instance, the heat capacity of (insulating) solids at low temperature is calculated by treating the solid as a gas of phonons.) But there is no way to identify a quantized vibration with any atom or collection of atoms in the solid – the best low-level translation of "there are such-and-such phonons present" would be "the solid is vibrating in such-and-such a way"

I agree with Wallace that the Nagelian way of doing reduction leads us astray, he claims that one goes wrong by attempting to translate the mathematical locution of physics into logic or natural language. The case of phonons also brings to light the idea that the world just isn't connected up in the way that the Nagelian assumes. Considering Nagel reduction to be a regulative ideal promotes the idea that one can find terms in the higher and lower-level models/theories that can be identified or at least regularly connected. But there is nothing at the atomic level that can be associated with a phonon – phonons are quantised modes of vibration of the lattice. The parts of the world all come together to realise the higher level in such a way that defies the piece-wise identification that the Nagel approach presumes. The derivation of phonons from the atomic description (as set out in part in Franklin and Knox (2018) following Kantorovich (2004)) is a good derivation and offers a sound reduction. However, it's not a Nagel reduction because there are no bridge laws.

Wilson (2006, p. 171) labels the idea that the world isn't connected up in such a way as to allow type-type identities 'distributed normativity'. He discusses a device that can be used to calculate the logarithm of a number and argues that no individual parts of the log function can be identified with parts of the device:

> mechanisms often illustrate a characteristic I shall call distributed normativity: some salient notion of "correctness" can be derived from the global purpose the device addresses. Consider, for example, the mechanical linkage ... whose purpose is to mechanically calculate the natural logarithm ($\ln(x)$) of the number selected by its left hand stylus. As such, the gizmo

> might prove useful in equilibrating the ratio of steam to fuel flow within an engine. This global ambition of calculating $\ln(x)$ naturally induces an internal evaluation of the "correctness" of the device's component parts–viz., have they all been sized properly to allow the complete mechanism to calculate $\ln(x)$ as ably as possible? I call such standards of "correctness" distributed because they filter down to the components of the mechanism from its overall purpose.

By analogy, the distributed support for higher-level terms suggests that constructing bridge laws will be challenging. Isolating physically meaningful terms in each part of each theory is not a trivial process; one often cannot determine what is connected to what. This means that connectability, which requires the identification of terms between lower and higher theories, may just not be possible. In the context of phonons this would amount to the claim that at the lower level there are particles arranged in a crystal lattice. But the individual phonons do not correspond to any particular set or arrangement of these particles. Rather once organised in this structure, and so long as they satisfy the harmonic approximation, the collective vibrations of these particles can be quantised and realise phonons. The Fourier transforms involved in phonon derivation are especially well suited to cases of distributed normativity and where they feature we ought to be extra suspicious of bridge laws in the service of reduction.

In this spirit Wilson (2021, p. 71) approvingly quotes Cassirer: "Physics is not a machine that can be taken apart. One cannot test every piece individually and wait until it has stood this sort of testing before putting it into the system. Physical science is a system that must be accepted as a self-contained whole, an organism of which no single part can be made to work without all the others, even those farthest removed." (Cassirer (1950, pp. 113–4)).

A further issue with the focus on Nagelian reduction as a model for reduction, and on formal methods at the expense of conceptual analysis, is that they can miss crucial problems. In the next section I take up the issues posed by my second case study which concerns the quantum-classical limit. I show that there that focussing on technical procedures, especially the $\hbar$ limit for relating classical to quantum physics, misses some relevant issues and in particular obscures the quantum measurement problem.

I raise this here in order to highlight one further aspect of the discussion of phase transitions that's been obscured at least in part by the literature's focus on the Nagel model, bridge laws and the justifiability of certain approximations/idealisations. What's assumed by both those who think that Nagelian reduction fails and those who argue that it succeeds is that the reason for that failure/success is in some sense intrinsic to the problem itself. The Nagelian reductionists argue that the infinite limit is required because of the limited tractability/computability of finite models but that better models could improve the situation and that there is no in principle issue at stake. The anti-reductionists tend to view this as an instance of the fact that the small scale and larger scale models of reality don't mesh. They regard the assumption that fundamental science is in-principle complete to be unwarranted and they view the failure of their preferred strict version of Nagelian reduction to be symptomatic of the broader failure of fundamental physics to do more than provide accurate models of the domains in which it has in fact been tested.

Wallace (2018) asks whether the problem of phase transitions might rather be just one more instance of the limits of applicability of closed system quantum physics. Thus it ought not to be dismissed as a non-problem. Systems that are small enough or effectively isolated from their environments sufficiently to maintain coherent states operate according to unitary quantum physics. If we were to assume that phase transitions are fundamentally governed by quantum physics (and thus break free of the classical assumptions in much of this literature) then we might notice that there is a very close resemblance between the issues considered above and those of the measurement problem. In particular one might view the symmetry-breaking evolution from one phase to another as the evolution from a coherent state to multiple decohered states, with a branching history on the decoherent histories formalism. This would explain how distinct outcomes with fewer symmetries can evolve from an initially more symmetric state.

When understood in terms of symmetry breaking the problem of phase transitions is: how does a system go from being modelled by the rotationally symmetric gas with free molecules to a system with fewer symmetries that's constrained in the way that liquids are? Or how is it possible for the system to choose between the ferromagnetic and antiferromagnetic conformation if the system is in the perfectly symmetric paramagnetic state? Wallace (2018) works through the suggestion that decoherence and an accompanying resolution of the measurement problem might be sufficient to resolve the issue; see also ideas discussed in Ruetsche (2011).

It's familiar that the measurement problem involves issues of symmetry breaking. So one might wonder why this suggestion of Wallace's – that there's a deep connection between problems of symmmetry breaking, including that of phase transitions, and the measurement problem – were not noticed before his intervention and have not been seriously taken up since.

I'd suggest that an important reason for this lacuna in the literature is that this approach doesn't fit with either side of the Nagelian model. Insofar as we rely on formal approaches we are expected to answer such questions either with the claim that phase transitions are reducible by derivation or that such a derivation fails and therefore phase transitions are emergent/autonomous/irreducible to lower-level descriptions. Reconceptualising in terms of the measurement problem transforms this issue but doesn't accord with the views of either camp.

In short, I think that there's grounds for optimism with respect to Wallace's analysis and that it's been missed precisely because attempting to answer questions of reduction rigorously with respect to a formal model serves to misdirect the debate.

The moral of this case is that unjustified idealisations can and do deserve philosophical reflection, and that Batterman and others who drew attention to these were right to do so, but that the bridge law/Nagelian approach can lead to the improper and unfair dismissal of such honest and well-justified philosophical exercises. Instead we should rather think hard about how the physics works and whether such idealisations can in fact be justified.

3 Case II: Nickles Reduction and the $\hbar$ Limit

The literature on reduction in the philosophy of physics, while more commonly focussed on Nagel reduction ('reduction$_1$') has a strand in which what's sometimes known as 'Nickles reduction' ('reduction$_2$') is cited as an alternative:

> In a broader sense, of course, reduction$_2$ is also "derivational". "Mathematical derivation", as the phrase is commonly used, includes not only logical deduction but limit processes and approximations of many kinds. . . . By contrast, reduction$_2$

> does not involve the theoretical explanation of one theory by another. ... the main functions of reduction$_2$ are justificatory and heuristic. The development of new theoretical ideas is heuristically guided by the requirement that these ideas yield certain established results as a special case (e.g., in the limit), and they are often quickly justified to a degree by showing that they bear a certain relation to a predecessor theory. ... The justification derives from the fact that the reduction shows the successor theory to account adequately for the structured domain of phenomena inherited from its successful predecessor
>
> [Nickles (1973, pp. 189, 185)]

Nickles, here, is after something less formal, as he terms it the relations between theories should be "justificatory and heuristic". Batterman takes this up and suggests that this is an alternative form of reduction, and that there are contexts in which the use of limits may provide reductive explanations. However, in Palacios (2022, pp. 556) we find the elevation of Nickles's views into another formal model of reduction: "A special case of Nickles' reduction is "limiting reduction," which refers to cases in which the transformations consist in mathematical limits. In cases where one limit is used, one can characterize *limiting reduction* as follows:

Limiting reduction: Let Q^1 denote a relevant quantity of T_1, Q^2 a relevant quantity of T_2, then a quantity Q^2 of T_2 $reduces_{lim}$ to a corresponding quantity Q^1 of T_1 iff (ii) $lim_{N\to\infty}Q^1_N = Q_2$ or $lim_{N\to 0}Q^1_N = Q_2$ (where N represents a parameter appearing in T_1) and (ii) the limiting operation makes physical sense."

While Palacios's formalisation of this approach adds rigour it's in line with much of Batterman's work in this area. As noted above he raises many important and crucial questions for putative reductions in philosophy of science. However, his argument that the nature of the limit matters, and that the singularity of the limit is a clear sign of reductive failure seems to take a model of reduction of this form too seriously and I fear that the extra precision Palacios provides increases the likelihood that one will make mistakes along these lines. The goal of this section is to demonstrate that appeal to the Nickles model as a paradigm for reduction impedes philosophical progress in discussions of the quantum-classical transition.

Questions of reduction in this context have been subject to a long debate, but what's especially relevant to us are the arguments from Batterman (2002) and Berry (1995) that the quantum-classical reduction must fail because the $\hbar \to 0$ limit is singular. Rosaler (2015) responds to this worry in line with the thoughts developed here: he argues that the focus on formal (theory-theory) reduction fails to observe that reduction should in fact be understood as an 'empirical' three part relation between the two theories to be inter-related and the world. Formal approaches to reduction leave out the idea that we can justify such limiting relations with reference to the world. Rosaler goes on to develop a detailed account of reduction where decoherence theory plays a crucial role in justifying the applicability of the $\hbar$ limit.

Yet Rosaler's analysis is rejected by Feintzeig (2022) in favour of a return to Nickles reduction and a focus on the limit as the primary relation to understand the reduction between classical and quantum physics. Feintzeig responds to Rosaler and claims that reduction in the formal sense is possible; he then provides formal reductions via C-star algebraic mathematical relations. He goes on to argue that empirical reduction à la Rosaler is inadequate to the scientific realist because it only offers the demonstration of continuity at the empirical level without the corresponding continuity at the formal level. Much more can be said about this particular debate and how to understand decoherence as providing a reduction that explains the continuity between the theories.

For our purposes I will restrict focus to a critique of formal reduction as provided by Feintzeig. In my view, the focus on formal reduction and limiting relations on the Nickles model is a distraction and does not provide the explanations of classical from quantum physics. I conclude by claiming that the literature on reduction would do much better to focus on explanation and understanding rather than the satisfaction of the criteria of any formal model. The upshot is that this will demonstrate the independence of questions of reduction from discussions of the singularity of the limiting relations.

Feintzeig (2022, p. 45) construes his work as a reductive project: "[w]e will take for granted that an intertheoretic reduction, if one is possible, would be constituted by an explanation of the success of classical mechanics on the basis of quantum mechanics". He goes on to claim that reductions will answer the question "why – if the temporal evolution of a physical system is better represented by quantum dynamics – can one (under some circumstances) approximately recover predictions for how expectation values and probability values change in time from classical dynamics?" (p.46).

Feintzeig goes on to detail this reduction and associates it directly with the limiting relation characterised above as Nickles reduction:

> Strict quantization treats the classical-quantum correspondence via the $\hbar \to 0$ limit, understood as a genuine limiting relation where $\hbar$ takes on numerical values and the algebraic relations between physical quantities vary in an appropriately continuous way between a classical and a quantum theory. ... The conditions of a strict quantization capture the idea that in the limit $\hbar \to 0$, the algebraic operations of the C*-product and commutator Lie bracket from the quantum theory approach the pointwise product and Poisson bracket of the classical theory
> [Feintzeig (2022, pp. 43, 27)]

I do think that there's value in the approach developed by Feintzeig, in particular it may underwrite a kind of structural continuity, also discussed by Thébault (2016) in the same context. In addition, as Feintzeig suggests it may help to understand how to quantise gravity. Nonetheless it's difficult to see how this approach can provide the kind of reductive explanation for which Feintzeig aims. This is because his approach cannot tell us under which circumstances classicality in fact emerges. If we are instructed to expect the equations of classical physics to be instantiated in the limit of $\hbar = 0$ we might remember that that limit is never in fact achieved. Assuming that the limit is a proxy for the ratio between the action and $\hbar$, that ratio is small for relatively large energies and long timescales, but on how long timescales and for how large energies does this hold?

What makes these questions more pressing is that $\hbar \to 0$ is not sufficient for classicality: there are increasingly many large systems that satisfy quantum descriptions in controlled experimental settings; see e.g. Fein et al. (2019) which reports interference of molecules with masses beyond 25,000 Da. It's standard to assume that explanations should be sufficient for their *explananda*! As such the $\hbar$ limit does not tell us why classical behaviour is found where it is, it does not explain the difference between contexts where we can reliably expect to find it and contexts where we cannot.

Feintzeig notes that he's after a "a *fixed error bound* such that the classical predictions lie within that error bound from the quantum predictions *no* *matter what state one began with*." It's difficult to see why one would expect that to be achievable given that systems instantiate different sets of equations in different contexts and neither classical nor quantum physics has a domain that can be specified in such abstraction. The $\hbar$ limit is far too crude a tool to allow for context-dependent description. And it's certainly subtle and context dependent where we can or should use classical vs. quantum physics.

Decoherence theory provides a detailed account of how quantum effects are screened off in certain conditions: insofar as interference terms are heavily suppressed we can derive that classical behaviour will be exhibited. Yet this account is context-dependent and non-rigorous in the sense that it's sensitive to the details of the models in each case and does not seem amenable to the kinds of mathematical physics that Feintzeig develops. Decoherence theory is a precise science that describes the applicability of classical physics to systems as they get larger, and it shows that smallness of $\hbar$ relative to action of the system is insufficient for classicality. What we need is precise details of how and to what extent interference is suppressed, only that can tell us when to expect classical behaviour.

What has gone wrong? Why has so much time been devoted to the rigorous and systematic study of the $\hbar$ limit without a clear articulation of what advantage it brings? In part this is answered by recent work by Landsman on the foundations of mathematical physics Landsman (2025): he emphasises that the strategy is to work on problems that are tractable. This might well be seen as an instance of Wilson's Physics avoidance (Wilson (2017)). The messy details on the way to the limit as $\hbar \to 0$ require case-by-case engagement, and so can conveniently be black-boxed by those interested in generality when considering the classical limit of quantum theory. But for the purposes of reduction this is inadequate.

What reduction requires is the intermediate ground between abstraction and universality, and this requires context sensitivity. Reduction doesn't stand or fall with anything as general as a limiting relation between theories nor does it depend on whether that limit is or is not singular. Decoherence theory is an exemplar of what's required for reduction: case by case analysis that pays attention to the details of how interference is supressed for each physical system under discussion. If full universality is required one can follow Feintzeig and attain a high level of rigour. But if one aims at a full explanation of where classicality is instantiated in fundamentally quantum systems, attention to the details of each and every decoherence model is required, and there will be a bespoke model for numerous different physical situations. This kind of analysis will not correspond to any formal model of reduction: the invocation of the formal models is incompatible with the detail necessary for the closing of explanatory gaps.

Limits play an important mathematical role in revealing the structure of any theory, thus they play an essential role in approximations and idealisations – they allow for abstractions that reveal details and they play methodological roles given the importance of approximately recovering certain equations in the limit. And they can thus serve explanations of the emergence and reduction between theories: they tell us where we should expect theoretical relations from one theory to be instantiated in models of another. However, they cannot be relied on to provide explanations and reductions because we never arrive at the limit, and so they require a substantial additional account of how it is that the limit reveals behaviour relevant to any explanatory story. They also can mislead in two ways, first they mislead by drawing focus to artefacts of the limits themselves, such as the singular nature of limits. Second they mislead by abstraction from details that matter: insofar as the decoherence functional is exactly zero and there is no interference between branches then any system evolves into just one branch and no quantum mechanical interpretation is needed. It's the approximate suppression of interference that highlights the sense in which decoherence doesn't solve the measurement problem; see Bacciagaluppi (2017).

4 Conclusion

This paper runs the risk of seeming polemical. Yet I do not wish to disparage the work of any of the philosophers raised in the above. Rather I think that the interesting work that they've done in various contexts has been misconstrued and many are thus misled by insistence that reduction should proceed according to one or other formal model of reduction. I claim that this particular set of issues in the philosophical analysis of inter-theoretic relations are instances of a much more general problem in philosophy of science and metaphysics: that an inappropriate standard of rigour serves as a regulative ideal.

If all else is equal rigour is surely a virtue for philosophical theorising. However, there's a trade-off between rigour and universality – one can, in principle, describe arbitrarily specific contexts with absolute rigour. But to describe some phenomenon or pattern that's multiply instantiated – that we find in different contexts with some variation in relevant detail – approximations and idealisations may defy the possibility of perfect rigour. The description of inter-theoretic relations are generically incompatible with high degrees of rigour for it's where different theories meet that science often fails to fill in all the details. This phenomenon is described at length in Wilson (2017) with the label 'physics avoidance'. Non-rigorous mathematical tools allow for approximation that can provide a much more detailed and complete explanation of the instantiation of one set of equations in contexts where others apply at more fundamental levels.

The two models of reduction and areas of the philosophical literature discussed in this paper should serve as cautionary tales against the idea that adding rigour is always virtuous, especially where this involves replacing approximative but successful scientific modelling with something more precise but lacking empirical success. More generally adding rigour may lead to clarity about standards for success in some contexts but undermine applicability to real scientific cases. The world just doesn't seem to be amenable to description by perfectly rigorous models, most clearly in contexts of inter-theoretic relations where the same modelling practices are expected to traverse domains. Physics uses whatever maths works in a given situation, whether or not this can be cleaned up and tidied is not obvious from the outset.

In slogan form: rigour is available at the extremes of generality and specificity but the explanations involved in reduction form an intermediate class, in which rigour is elusive. Nickles and Nagel reduction may be available and informative in some contexts but neither serves as a good regulative ideal for reduction, and may underwrite the misapprehension that inter-theoretic relations are simpler than they are! Mathematical derivations serve an important function, as do formal approaches, but one should not assume that the philosophical work stops there, nor should one attempt to hammer these into a preferred framework.

The conclusion is not 'don't be rigorous!', but rather: 'think hard about where rigour is justified and warranted and where it isn't'.

Acknowledgements

Thanks to the audience at the workshop on Rigour in Paris, and the Metaphysics and Science reading group at King's College London, as well as Karim Thébault and Jingyi Wu for extended discussions on these issues.

References

Bacciagaluppi, Guido (2017). "The Role of Decoherence in Quantum Mechanics". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Summer 2017. Metaphysics Research Lab, Stanford University.

Bangu, Sorin (2009). "Understanding Thermodynamic Singularities: Phase Transitions, Data, and Phenomena*". In: *Philosophy of Science* 76.4, pp. 488–505.

— (2015). "Why Does Water Boil? Fictions in Scientific Explanation". In: *Recent Developments in the Philosophy of Science: EPSA13 Helsinki*. Ed. by Uskali Mäki et al. Cham: Springer International Publishing, pp. 319–330.

Batterman, Robert W. (Sept. 1992). "Explanatory Instability". In: *Noûs* 26.3, pp. 325–348.

— (2000). "Multiple Realizability and Universality". In: *The British Journal for the Philosophy of Science* 51.1, pp. 115–145.

Batterman, Robert W. (2002). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press.

— (2005). "Critical phenomena and breaking drops: Infinite idealizations in physics". In: *Studies In History and Philosophy of Modern Physics* 36.2, pp. 225–244.

— (2010a). "On the Explanatory Role of Mathematics in Empirical Science". In: *The British Journal for the Philosophy of Science* 61, pp. 1–25.

— (2010b). "Reduction and renormalization". In: *Time, Chance, and Reduction: Philosophical Aspects of Statistical Mechanics*. Ed. by Gerhard Ernst and Andreas Hüttemann. Cambridge University Press, pp. 159–179.

— (2011). "Emergence, singularities, and symmetry breaking". In: *Foundations of Physics* 41, pp. 1031–1050.

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

Batterman, Robert W. and Collin C. Rice (2014). “Minimal model explanations”. In: *Philosophy of Science* 81.3, pp. 349–376.

Berry, Michael (1995). “Asymptotics, singularities and the reduction of theories”. In: *Studies in Logic and the Foundations of Mathematics*. Vol. 134. Elsevier, pp. 597–607.

Bokulich, Alisa (2008). *Reexamining the Quantum-Classical Relation: Beyond Reductionism and Pluralism*. Cambridge University Press.

Butterfield, Jeremy (June 2011). “Less is Different: Emergence and Reduction Reconciled”. In: *Foundations of Physics* 41.6, pp. 1065–1135.

Butterfield, Jeremy and Nazim Bouatta (2012). “Emergence and Reduction Combined in Phase Transitions”. In: *AIP Conference Proceedings* 1446, pp. 383–403.

Callender, Craig (2001). “Taking thermodynamics too seriously”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 32.4, pp. 539–553.

Cassirer, Ernst (1950). *The Problem of Knowledge*. New Haven: Yale University Press.

Craver, Carl F. and David M. Kaplan (2018). “Are More Details Better? On the Norms of Completeness for Mechanistic Explanations”. In: *The British Journal for the Philosophy of Science*.

Dizadji-Bahmani, Foad (2021). “Nagelian reduction in physics”. In: *The Routledge Companion to Philosophy of Physics*. Routledge, pp. 499–511.

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

Fein, Yaakov Y. et al. (Dec. 2019). “Quantum superposition of molecules beyond 25 kDa”. In: *Nature Physics* 15.12, pp. 1242–1245. ISSN: 1745-2481. DOI: 10.1038/s41567-019-0663-9. URL: https://doi.org/10.1038/s41567-019-0663-9.

Feintzeig, Benjamin H (2022). *The classical–quantum correspondence*. Cambridge University Press.

Franklin, Alexander (2021). “Can Multiple Realisation Be Explained?” In: *Philosophy* 96.1, pp. 27–48.

— (2024). “How the Reductionist Should Respond to the Multiscale Argument, and What This Tells Us About Levels”. In: *Levels of Explanation*. Ed. by Katie Robertson and Alastair Wilson. Oxford University Press. Chap. 3. DOI: 10.1093/oso/9780192862945.001.0001.

Franklin, Alexander and Eleanor Knox (2018). “Emergence without limits: The case of phonons”. In: *Studies In History and Philosophy of Modern Physics* 64, pp. 68–78.

Kadanoff, Leo P. (2009). “More is the Same; Phase Transitions and Mean Field Theories”. In: *Journal of Statistical Physics* 137, pp. 777–797.
Kantorovich, Lev (2004). *Quantum theory of the solid state: an introduction*. Vol. 136. Springer Science & Business Media.
Knox, Eleanor (2016). “Abstraction and its Limits: Finding Space For Novel Explanation”. In: *Noûs* 50.1, pp. 41–60.
Landsman, Klaas (June 2025). *Philosophy of Mathematical Physics*. URL: https://philsci-archive.pitt.edu/25843/.
Liu, Chuang (1999). “Explaining the Emergence of Cooperative Phenomena”. In: *Philosophy of Science* 66, S92–S106.
Mainwood, Paul (2006). “Is More Different? Emergent Properties in Physics”. PhD thesis. University of Oxford.
Morrison, Margaret (2006). “Emergence, Reduction, and Theoretical Principles: Rethinking Fundamentalism”. In: *Philosophy of Science* 73.5, pp. 876–887.
Nagel, Ernest (1961). *The Structure of Science: Problems in the Logic of Scientific Explanation*. Second Edition (1979). Hackett Publishing.
Nguyen, C. Thi (June 2020). “Gamification and Value Capture”. In: *Games: Agency As Art*. Oxford University Press. DOI: 10.1093/oso/9780190052089.003.0009.
Nickles, Thomas (1973). “Two concepts of intertheoretic reduction”. In: *The Journal of Philosophy* 70.7, pp. 181–201.
Palacios, Patricia (2022). *Emergence and reduction in physics*. Cambridge University Press.
Potochnik, Angela (2017). *Idealization and the Aims of Science*. University of Chicago Press.
Putnam, Hilary (1975). “Philosophy and our mental life”. In: *Mind, Language and Reality: Philosophical Papers, Volume 2*. Cambridge University Press, pp. 291–303.
Railton, Peter (1981). “Probability, Explanation and Information”. In: *Synthese* 48.2.
Robertson, Katie (2019). “Reductive Aspects of Thermal Physics”. PhD thesis. University of Cambridge.
Rosaler, Joshua (2015). ““Formal” versus “Empirical” Approaches to Quantum–Classical Reduction”. In: *Topoi* 34.2, pp. 325–338.
Ruetsche, Laura (2011). *Interpreting Quantum Theories: The Art of the Possible*. Oxford University Press.
Schaffner, Kenneth F. (2013). “Ernest Nagel and Reduction”. In: *The Journal of Philosophy* 109.8/9, pp. 534–565.

Styer, Daniel F. (2004). "What good is the thermodynamic limit?" In: *American Journal of Physics* 72.1, pp. 25–29.

Thébault, Karim PY (2016). "Quantization as a guide to ontic structure". In: *The British Journal for the Philosophy of Science*.

Wallace, David (2018). *Spontaneous Symmetry Breaking in Finite Quantum Systems: a decoherent-histories approach*. arXiv: `1808.09547 [quant-ph]`.

— (2022). "Stating structural realism: mathematics-first approaches to physics and metaphysics". In: *Philosophical Perspectives* 36.1, pp. 345–378.

Wilson, Mark (2006). *Wandering Significance: An Essay on Conceptual Behaviour*. Oxford University Press.

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

— (2021). *Imitation of rigor: An alternative history of analytic philosophy*. Oxford University Press.

Yablo, Stephen (1992). "Mental causation". In: *The Philosophical Review* 101.2, pp. 245–280.