Philosophy of Science (2026), 1–10
doi:10.1017/psa.2026.xxxxx

PSA PHILOSOPHY OF SCIENCE ASSOCIATION PHILOSOPHY OF SCIENCE

CONTRIBUTED PAPER

The Structure of Phonons

Cameron Dashwood

Department of History and Philosophy of Science, University of Cambridge, Cambridge, UK
Email: cdd39@cam.ac.uk

(Received June 1, 2026)

Abstract
Phonons—particle-like quanta that emerge from collective vibrations of the atoms that make up materials—have been touted as promoting a structuralist ontology. Structural realists claim to avoid the difficulties of bridging the ontological gap between phonons and the underlying atoms by committing only to the mathematical structure shared by their theories. I scrutinise this claim by analysing the shared structure and considering what motivates realism about it. I identify a tension between the demands of structural continuity and experimental testing, and consider the viability of supplementing the structuralist ontology with some form of operational structure.

1. Introduction

Phonons are everywhere. They are the carriers of heat and sound in solids (and some liquids). They produce electrical resistance in metals by deflecting conduction electrons from their path. And, in superconductors, they even bind electrons together into *Cooper pairs* that flow without resistance. They are formally analogous to photons: just as photons are quanta of light, so phonons are quanta of sound. Both are described by quantum field theories, have particle-like properties of energy and momentum, and obey Bose-Einstein statistics.

Given their centrality to many theoretical explanations in condensed matter physics, phonons would seem a good candidate for realism. Indeed, a survey of realist attitudes among physicists found strong agreement with the statement "phonons exist" (Henne et al., 2024). The agreement from philosophers was significantly lower, however, especially when compared to the statement "electrons exist". This response likely stems from a key difference between phonons and electrons: whereas electrons are fundamental particles, phonons are emergent *quasiparticles* that arise from the collective vibrations of atoms. Gelfert (2003, 246) voices the sentiment that quasiparticles "are no proper kinds of entities at all—they are merely collective effects of (an indeterminate number of) 'real' entities, and they must be acknowledged as *illusory*".

The problem for the would-be realist about phonons is finding a way to extend their realist commitment to phonons without generating ontological inflation and disunity. What is needed is some way of bridging the ontologies of phonons and the underlying atoms. Unfortunately, traditional mereological approaches provide little help. As Gelfert (2003, 261) notes, the relation between phonons and atoms "is one of correlation rather than containment". A phonon cannot be identified with any particular collection of atoms, and each atom simultaneously contributes to a variable number of phonons (and other types of quasiparticle). Classically, each phonon can be understood as corresponding to a pattern of atomic vibrations, but even this picture breaks down at the quantum level where displacements become indeterminate.

Structural realism offers an inviting alternative approach. It identifies the root of the traditional realist's predicament as their adherence to an object-based ontology. We should not take scientific theories to be telling us about objects like phonons and atoms, says the structural realist, but about the *structure* of the world. This structure is claimed to show greater continuity on theory change, and is therefore a more robust thing to hang our realism on. This perspective on phonons has been particularly advocated by adherents of *ontic structural realism* (OSR) (Ladyman, 2018; Wallace, 2003, 2022). Rather than an epistemic constraint on what we can know, OSR is a metaphysical thesis about how the world is (Ladyman, 1998). Motivated by issues around the individuality of quantum particles (like phonons) and the highly mathematical nature of modern physical theories (like phonon theory), OSR takes the mathematical structure of theories to be ontologically basic.

Despite appealing to phonons to support their position, structural realists have not offered an account of the structure shared by the atomic and phonon theories, nor established that standard arguments for realism apply to it. In this paper, I will attempt to fill these gaps, and in so doing to scrutinise the structuralist perspective on phonons. I will first introduce the phonon theory through a toy model of a chain of vibrating atoms (section 2.1). Having used this model to identify the common structure of the atomic and phonon theories (section 2.2), I will then consider what motivates realism about it. I argue that operationalisation is needed to connect theories to phenomena (section 3.1), and that operationalisation requires some physical interpretation of structural parameters (section 3.2). I find, however, that the interpretations are not preserved between the theories. I conclude by considering whether OSR can escape this predicament through some form of 'operational structure' (section 4).

2. The theory of phonons

The history of phonons goes back to the origins of the old quantum theory and follows the development of both quantum mechanics (QM) and quantum field theory (QFT). This history has yielded a rich set of theoretical tools for modelling phonons and their influence on the properties of a range of materials. Luckily, the relation between phonons and the underlying atoms can be captured by a simple toy model of a one-dimensional chain. I will briefly describe this model in the next section, including some illustrative equations but omitting as much detail as possible and hoping that the ideas are clear even if the equations are opaque. Further details can be found in condensed matter textbooks (such as Altland and Simons, 2006).

2.1. Phonons in a one-dimensional chain

Here is the setup. Imagine a one-dimensional chain of N identical quantum-mechanical atoms of mass m, with equilibrium spacing d, and with harmonic interactions between neighbouring atoms. The behaviour of this chain is dictated by the *Hamiltonian*—the total energy operator—which in this case is just the sum of the kinetic and potential energies:

$$\hat{H} = \sum_{i=1}^{N} \frac{\hat{p}_i^2}{2m} + \sum_{\langle i,j \rangle} \frac{\lambda}{2} \left(\hat{x}_i - \hat{x}_j\right)^2 , \tag{1}$$

where $\hat{x}_i$ and $\hat{p}_i$ are respectively the position and momentum operators of the i^{th} atom, λ is the interaction strength, and $\langle i, j \rangle$ indicates that the second sum is taken only over neighbours.

Technically, equation 1 (plus some boundary conditions) is all we need to calculate the evolution of the system via the Schrödinger equation. For this simple system, there are some tricks that can be used to solve the Schrödinger equation by hand. But for systems involving more complex atomic structures and non-harmonic potentials, numerical methods are required that become intractable for values of N anywhere near the $\sim 10^{23}$ of macroscopic materials. Physicists interested in the large-scale properties of such systems can turn the high N from a problem to an advantage, however, by 'zooming out' from the details of the atoms, replacing their description with one of continuous fields and employing the techniques of QFT.

Moving from atomic to field-theoretic descriptions involves taking the continuum limit, in which the number of atoms goes to infinity ($N \to \infty$) and the spacing to zero ($d \to 0$) while preserving the total length ($L = Nd$). Classically, this corresponds to a discrete chain becoming a continuous solid, with displacements of the atoms becoming density variations. Formally, it involves replacing the set of displacement operators, $\{\hat{x}_i\}$, with an operator-valued field, $\hat{\phi}(x)$, over a continuous variable x (note x is no longer an operator but a coordinate that replaces the indices i). The Hamiltonian is then written in terms of this quantum field as

$$\hat{H} = \int_0^L \left[\frac{m}{2} \left(\frac{\partial \hat{\phi}}{\partial t} \right)^2 + \frac{\lambda d^2}{2} \left(\frac{\partial \hat{\phi}}{\partial x} \right)^2 \right] dx. \tag{2}$$

Following the canonical quantisation approach, we can then transform to the *Fock representation*, in which the field is described as a collection of *quanta*—packets of energy—which in this case are phonons. The Hamiltonian becomes

$$\hat{H} = \sum_k \hbar \omega_k \left(\hat{a}_k^\dagger \hat{a}_k + \frac{1}{2} \right) , \tag{3}$$

where the space now is indexed by the *wavevector* $k = 2\pi n/L$ with $n \in \mathbb{Z}$ (the Fourier inverse of x). Those familiar with QFT will recognise this Hamiltonian as being identical to that of the free electromagnetic field, where $\hat{a}_k^\dagger$ and $\hat{a}_k$ are called *creation* and *annihilation* operators because they respectively create or destroy a photon. Similarly, we can interpret $\hat{a}_k^\dagger/\hat{a}_k$ here as creating/annihilating a phonon with particle-like momentum $\hbar k$ and energy $\hbar\omega_k = \hbar d \sqrt{\lambda/m}|k|$. The *ground state*—the state of minimum energy—is the one containing zero phonons: the phonon *vacuum*. The combination $\hat{a}_k^\dagger \hat{a}_k$ can be shown to simply count the number of phonons of that k-value, with the Hamiltonian summing up these numbers multiplied by the respective energies (plus the *zero-point energy*, $\sum_k \hbar\omega_k/2$).

2.2. Bridging the ontological gap

I have introduced two theories—the atomic and phonon theories—related by an asymptotic limit. They have overlapping domains of application but describe very different ontologies. The former describes a discrete array of eternal quantum-mechanical atoms vibrating around their equilibrium positions. In the latter there are no atoms and no vibrations. Instead, there is a continuous field made up of spatially delocalised phononic quanta created out of a vacuum. How can we bridge the ontological gap between them?

We have already seen that object-based approaches like mereology are not the answer. OSR instead urges us to seek continuity at the level of structure. If a common structure can be identified, then we can be realist about that structure without having to worry about incommensurable objects. The task is then to determine the common structure. It is surprisingly difficult to find precise definitions of structure in the OSR literature, so I take my cue from Ladyman and Ross (2007, 159) describing the structure of general relativity (GR):

> According to OSR, if one were asked to present the ontology of the world according to, for example, GR one would present the apparatus of differential geometry and the field equations and then go on to explain the topology and other characteristics of the particular model (or more accurately equivalence class of diffeomorphic models) of these equations that is thought to describe the actual world.

We can glean from this that the structure of interest is the mathematical structure encoded in the equations that is invariant under the symmetries of the theory (this is what is meant by “equivalence class of diffeomorphic models”). Applying this to the atomic theory, the structure is something like the total *Hilbert space*—the tensor product of the individual spaces of complex square-integrable functions for each atom—with the state of the system represented by a *ray*—a vector with arbitrary complex phase—whose evolution is dictated by the Schrödinger equation containing the Hamiltonian operator of equation 1. This structure satisfies the features identified above, being invariant under spatiotemporal and permutation symmetries. It is also insensitive to the choice of Heisenberg or Schrödinger representations, as the Stone–von Neumann theorem ensures these have *isomorphic*—structurally equivalent—Hilbert spaces.

With the structure of the atomic theory in hand, we then need to find its overlap with the structure of the phonon theory. The intuitive method would be to pick out a subset of the degrees of freedom that can be identified with those of the phonon theory. Wallace (2022, 357) proposes just such a procedure: “the phonon description of a solid can be simply seen as instantiated by the atomic-level description just by showing a (spatial-symmetry-preserving) Hilbert-space isomorphism between the [theories]”. Although simple to state, this procedure is complicated by fact that the phonon theory has more degrees of freedom than the atomic theory. In fact, it has *infinite* degrees of freedom: the value of the field at each infinitesimal point. We must therefore first cut down the degrees of freedom of the phonon theory, for instance by introducing a spatial cutoff that discards dynamics at scales below d.

Carrying out this procedure would be tedious and unilluminating. What is more important for our purposes is the conclusion that not all the mathematical structure of the phonon theory survives in the atomic theory; structure is both lost and gained on theory change. This situation is hardly surprising, being bound up with the nature of the asymptotic limit. Nor is it unique to the phonon case: physics is replete with reductions of continuum by atomic/statistical theories. As Ladyman (2021) acknowledges, OSR is easily refuted if it requires that all mathematical structure is retained on theory change. Instead, OSR must fall back on the *approximate* preservation of structure—their version of traditional realists' retreat to 'approximate truth'.

The demands of realism do still place constraints on this approximate preservation, however. It would be problematic if the structure retained on one theory change was discarded under successive changes, for instance. OSR requires that there is some robust, accumulating core of structure. It also seems crucial that any discarded structure is inessential to the empirical success of the reduced theory. Otherwise, its loss undermines the inference from the empirical success of that theory to the reality of its structure. Determining whether the first of these conditions is met would require a survey of theories that are reduced by the phonon theory or reduce the atomic theory, and is beyond the scope of this paper. Instead, I will turn to the second condition. Having seen how OSR can (approximately) get around the main argument against realism, we now need to investigate whether it can satisfy the standard argument for realism.

3. Making contact with the empirical

The argument for structural *realism* is a variant of the no-miracles argument to the effect that the success of science would be a miracle unless its structures were latching onto reality (Worrall, 1989; Ladyman, 1998; Ladyman and Ross, 2007; Wallace, 2022). Whether this 'success' is cashed out as the ability to make novel predictions, unify disparate phenomena, or simply make highly accurate predictions, it is clear that the ability of OSR to fulfil the promise of the no-miracles argument depends on its account of how theories make contact with empirical phenomena.

Most OSR advocates subscribe to the *semantic view* in which theories are connected to phenomena through some kind of *morphism*—a mapping that preserves mathematical structure (French and Ladyman, 1999; Ladyman and Ross, 2007; Wallace, 2022). Different authors prefer different types of morphism, including *iso*-, *homo*-, and *partial*-morphisms. The technical details need not concern us. What is important is that they are all, in some sense, global mappings: they do not depend on any similarity between each element of the source and target beyond their structural relations to the other elements.

The morphism approach initially appears to work well in the phonon case. Indeed, we already met isomorphism in section 2.2 as the relation between different Hilbert-space representations and as a method to find the common structure of the atomic and phonon theories. Broadening this from inter-theory to theory–data relations, we can imagine establishing morphisms between theoretical structures and data models. For example, a commonly measured property of phonons is their *dispersion relation*—the dependence of their energy on momentum. Given a data set containing phonon energies and momenta for a particular material, we can generate a data model by fitting a curve through the points. We can also extend our theoretical model from section 2.1 to include the full atomic structure of the material. We can then establish a morphism between the data model and the substructure of the theoretical model that encodes the dependence of ω_k on k.

But this has only taken us so far. Morphisms might provide a structuralist account of inter-theory and theory–data relations, but, being defined between two mathematical structures, they are unable to make the further connection to raw phenomena. As Suárez (2003), Frigg (2006), and Brading and Landry (2006) have pointed out, phenomena do not present themselves to us pre-structured; work has to be done to structure them. French and Ladyman (1999, 112–113) concede this point:

> [P]henomena, understood presumably as something like actual physical events, cannot be put into *isomorphism* [...] since the notion of isomorphism is defined for mathematical objects only [...] Of course there is the more profound issue of the relationship between the lower most representation in the hierarchy—the data model perhaps—and reality itself, but of course this is hardly something that the semantic approach alone can be expected to address.

This phrasing indeed makes the relation between theories and 'reality' appear mysterious, which may be what motivates OSR to sidestep the problem by declaring that reality *just is* mathematical structure. Suárez, Frigg, and Brading and Landry, meanwhile, diagnose the semantic view as unable to provide a satisfactory account of scientific representation. But, as philosophers of experimentation going back to Hacking (1983) have taught us, we cannot make direct contact with phenomena if we stay in the realm of representation. To understand the theory–world relation, we need to enter the laboratory.

3.1. Testing structures in the lab

In experimental practice, the theory–phenomenon connection is established by the physical operations used to measure aspects of a phenomenon and assign numerical values to corresponding theoretical parameters. The focus on measurement operations finds its philosophical home in *operationalism*, a now-shunned position that arose from the writings of physicist Percy Bridgman. At his most extreme, Bridgman (1927, 5) held that "we mean by any concept nothing more than a set of operations; the concept is synonymous with the corresponding set of operations". This radical doctrine has been rightly criticised as fragmenting empirical concepts into the different ways of measuring them and evacuating theoretical concepts of meaning. I do not intend to endorse operational*ism* as a reductive theory of meaning, but follow Chang (2004) in taking operational*isation* as a useful framing of the process by which scientists connect theories to phenomena.

Operationalisation differs markedly from the construction of morphisms. Crucially, it is not a purely representational practice (although it can include mental—or what Bridgman called "paper-and-pencil"—operations, alongside physical manipulations). It also operates at a different structural level, applying not to whole structures but to elements within a structure. Recall that I characterised morphisms as global mappings, in the sense that they depend only on the relations of elements with the rest of the structure. Operationalisation is comparatively piecemeal. As Chang (2004, 197) states, "[o]perationalizing an abstract theory involves operationalizing certain individual concepts occurring in it, so that they can serve as clear and convenient bridges between the abstract and the concrete".

Take the story I told above about dispersion relations. To complete this story we need to ask how the data sets are generated in the first place. This requires operationalising the parameters ω_k and k, for instance through *inelastic x-ray scattering* (Burkel, 2000). Roughly, ω_k can be operationalised by: firing a high-intensity beam of x-ray photons at a sample; diffracting the scattered photons from a silicon crystal and measuring the diffraction angle, θ; calculating the wavelength of the scattered photons using Bragg's law, $\lambda = 2d \sin\theta$ (where d is the atomic spacing of silicon); and finally taking the difference from the wavelength of the incident photons and converting it to an energy via $\omega = 2\pi c/\lambda$. Simultaneously, k can be operationalised by: measuring the angle through which the photons are scattered; using Euclidean geometry to calculate the vector difference between the incident and scattered beams, each with magnitude $2\pi/\lambda$; and finally taking the magnitude of this difference vector. It is only once these operations are repeated for different values of ω_k and k that a data model can be generated and the morphism approach applied.

3.2. The need for interpretation

We have seen that operationalisation is needed to connect theories to phenomena. But what determines the operationalisation of a theoretical parameter? The best situation for the structural realist would be if operationalisation was determined by the structural relations of the parameter to the rest of the theory, such that the mathematical structures identified in section 2.2 dictated their own operationalisations.

Unfortunately, this situation is ruled out by the scattering process described above: the formal analogy between phonons and photons means the parameters ω_k and k are structurally equivalent in the theories, but they are not operationally equivalent. For phonons, these parameters can be operationalised by photon scattering, but the equivalent parameters of the photon theory cannot be operationalised by phonon scattering. This asymmetry is largely a consequence of the fact that we cannot get phonons 'out' of their material hosts. We can measure ω_k and k of a photon by diffracting it from a silicon crystal and detecting it using the photoelectric effect, but to measure these same properties of a phonon we have to go via a scattered photon.

This is one example of the wider issue that isomorphic structures are used to represent an array of disparate phenomena. Frigg and Votsis (2011, 261) mention it as a problem for OSR, giving as a further example the fact that "the elongation of a lead ball bouncing up and down on a metal spring and the voltage over a condenser obey the same equation and, in that sense, have the same formal structure". In fact, the disconnect between mathematical structure and operationalisation is even more stark in this case. Again, structurally equivalent parameters are operationalised differently: the 'amplitude' parameter might be operationalised with a ruler for the ball-on-a-spring system, versus a voltmeter for the condenser circuit. Not only that, but the same operationalisations can apply to parameters that occupy different positions in isomorphic structures: compare an oscillating circuit where the components are placed in series to one where they are placed in parallel.

So, if a theoretical structure cannot guide its own operationalisation, what can? It seems that what is needed is some *physical interpretation* of the parameters. Operationalising the amplitude parameter for the ball-on-a-spring, for instance, requires it to be interpreted as a length, such that pre-existing operationalisations of length involving rulers can be applied. I note that structural realists routinely partake in this kind of interpretation. For examples, Ladyman and Ross (2007, 95) express Ehrenfest's theorem as "grad$V(\langle r \rangle) = m(\mathrm{d}^2\langle r \rangle/\mathrm{d}t^2)$, where V is the potential and r is the position operator", while Wallace (2022, 348) describes the structure of Newtonian mechanics as containing "N positive real numbers $m_1, \ldots m_N$, representing the particle masses".

Returning to phonons, the operationalisations I gave for $\hbar\omega_k$ and $\hbar k$ appear quite mysterious unless these parameters are interpreted as an energy and a momentum. Then it can be understood that, as conserved quantities, they will be equal and opposite to the change in photon energy/momentum on scattering, and existing operationalisations of photon energy and momentum can be employed. It is not sufficient to interpret these parameters as an energy or momentum simpliciter, however, as we have seen that they are operationally distinct for phonons and photons. This suggests even finer distinctions are necessary—distinctions that would traditionally be understood as circumscribing the properties of different physical objects.

I will not attempt to delimit the extent of interpretation needed for operationalisation. Wherever the line is drawn (I suspect that it is neither sharp nor static), there is an immediate problem for OSR: interpretation opens the door to discontinuity on theory change. This is clearly the case if the parameters are interpreted as the properties of self-subsistent objects, as they are then exactly the things we discarded in section 2.2. But even the most minimal interpretations are liable to revision. Return to the interpretation of $\hbar k$ as a momentum (ignoring further distinctions about the type of momentum). From the perspective of the phonon theory, this is a true momentum that is conserved and exchanged with photons on scattering. We would then expect the total momentum of the phonon field, with operator

$$\hat{P}_{\text{phonon}} = -\int_0^L m \frac{\partial \hat{\phi}}{\partial t} \frac{\partial \hat{\phi}}{\partial x}\, dx = \sum_k \hbar k\, \hat{a}_k^\dagger \hat{a}_k, \tag{4}$$

to be structurally equivalent to the operator that gives the total momentum in the atomic theory, $\hat{P}_{\text{atomic}} = \sum_i \hat{p}_i$. But this is not the case: in the continuum limit, $\hat{P}_{\text{atomic}} \to \int m(\partial\hat{\phi}/\partial t)\,\mathrm{d}x \neq \hat{P}_{\text{phonon}}$.

We can define a parameter in the atomic theory that is the structural equivalent of k, given by $k_{\text{atomic}} = 2\pi n/L$ with $n = 0, \pm 1, \pm 2, \ldots \pm N/2$. But, despite the structural equivalence, $\hbar k_{\text{atomic}}$ is not interpreted as a momentum; instead it is a *quasi*-momentum that is not generally conserved. In fact, from the perspective of the atomic theory, all finite-energy excitations with $k_{\text{atomic}} \neq 0$ have *zero* total momentum. This necessitates a new understanding of photon scattering. Now, energy is understood as being transferred to zero-momentum excitations with $k_{\text{atomic}} \neq 0$, while momentum is transferred to additional zero-energy excitations with $k_{\text{atomic}} = 0$.

This interpretive revision is related to the loss of mathematical structure we encountered in section 2.2. In the phonon theory, the continuity of x results in an unbounded k whose conservation is a consequence of invariance under x-translations according to Noether's theorem. In the atomic theory, by contrast, the discreetness of i results in k_{atomic} being bounded to $[-\pi/d, \pi/d]$ and not conserved as Noether's theorem does not apply to discrete symmetries. Of course, being conserved is not a sufficient condition to interpret a parameter as a momentum, but it is necessary. The fact that mathematical structure is only approximately preserved therefore leads to important differences in our understanding of how the theories are tested by experiment.

We have ended up in a double bind. On the one hand, structural realists must undertake some interpretation of their mathematical structures in order to connect them to phenomena. This is needed to establish the empirical success that fuels the no-miracles argument and therefore underwrites their realism. On the other hand, to engage in even a minimal amount of interpretation reintroduces the kind of referential discontinuity whose avoidance is the raison d'être of OSR. Structural realists appear to be left to choose between their structuralism and their realism.

4. Operational structure?

At this point, structural realists might cry foul. They might admit that some interpretation is needed, but argue that to take it as consisting in empirical concepts like 'momentum' is to fall back into object-based ways of thinking (despite their own use of these concepts); instead, interpretation requires more structure. Ladyman and Ross (2007, 122) seem to get at something like this with their concept of "locators", each "a partial interpretation of a structure in the context of another, presupposed, structure", where by "partial interpretation" they mean "an indication of where in the universe measurements should be taken that will be relevant to assigning values to variable parameters of the structural element in question".

Operationalising a structure clearly requires more than locating its instantiations in spacetime. But can a more nuanced version of this tactic succeed? What is needed is some form of 'operational structure' that connects theoretical parameters to operationalisations without going via substantive empirical concepts. We could imagine these extra-theoretical structures as graphs, with the nodes on one side being the operationally equivalent parameters across theories, linked on the other side to their different operationalisations.

The crucial question is whether operational structures are better preserved than interpretations in terms of empirical concepts. The situation does look more hopeful. If one works through scattering theory applied to the atomic chain, it turns out that the value of k_{atomic} is related to the change in momentum of the scattered photon, despite not being a momentum itself. There is therefore some continuity in the operational structure of k and k_{atomic}. The loss of fine-scale mathematical structure on moving between the theories still leads to differences in the details of their operationalisations, however. Whereas the loss of photon momentum can be directly equated with the unbounded k, it must be projected into the range $[-\pi/d, \pi/d]$ before being equated with k_{atomic}.

The move from substantive interpretations to operational structure can therefore reduce discontinuity to a level commensurate with the loss of mathematical structure that OSR is already saddled with. However, it raises a number of additional issues. First, it should be remembered that the phonon case is particularly suited to a structuralist reading, and it seems likely that cases involving more significant shifts in mathematical structure will show more problematic discontinuities in operational structure. Second, the plasticity of empirical concepts appears to play a generative role in developing new operationalisations, and throwing them out may fossilise the products of science. Third and finally, it must be appreciated just how different mathematical and operational structure are. Whereas it is plausible that mathematical objects can be fully defined by their structural relations (Benacerraf, 1965), this is clearly not the case for physical operations. As well as reigniting questions of what relations mean without relata, subscribing to operational structure threatens to trap us again in the abstract realm of representation.

References

Altland, A. and Simons, B. D. (2006) *Condensed Matter Field Theory*. Cambridge University Press. Cambridge. https://doi.org/10.1017/CBO9780511789984.

Benacerraf, P. (1965) What numbers could not be. *The Philosophical Review*. *74*, 47–73. https://doi.org/10.2307/2183530.

Brading, K. and Landry, E. (2006) Scientific structuralism: Presentation and representation. *Philosophy of Science*. *73*, 571–581. https://doi.org/10.1086/518327.

Bridgman, P. W. (1927) *The logic of modern physics*. Macmillan. New York. https://archive.org/details/logicofmodernphy00brid.

Burkel, E. (2000) Phonon spectroscopy by inelastic x-ray scattering. *Reports on Progress in Physics*. *63*, 171. https://doi.org/10.1088/0034-4885/63/2/203.

Chang, H. (2004) *Inventing Temperature: Measurement and Scientific Progress*. Oxford University Press. Oxford. https://doi.org/10.1093/0195171276.001.0001.

French, S. and Ladyman, J. (1999) Reinflating the semantic approach. *International Studies in the Philosophy of Science*. *13*, 103–121. https://doi.org/10.1080/02698599908573612.

Frigg, R. (2006) Scientific representation and the semantic view of theories. *Theoria*. *21*, 49–65. https://doi.org/10.1387/theoria.553.

Frigg, R. and Votsis, I. (2011) Everything you always wanted to know about structural realism but were afraid to ask. *European Journal for Philosophy of Science*. *1*, 227–276. https://doi.org/10.1007/s13194-011-0025-7.

Gelfert, A. (2003) Manipulative success and the unreal. *International Studies in the Philosophy of Science*. *17*, 245–263. https://doi.org/10.1080/0269859032000169451.

Hacking, I. (1983) *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge University Press. Cambridge. https://doi.org/10.1017/CBO9780511814563.

Henne, C., Tomczyk, H. and Sperber, C. (2024) Physicists' views on scientific realism. *European Journal for Philosophy of Science*. *14*, 10. https://doi.org/10.1007/s13194-024-00570-z.

Ladyman, J. (1998) What is structural realism? *Studies in History and Philosophy of Science*. *29*, 409–424. https://doi.org/10.1016/S0039-3681(98)80129-5.

Ladyman, J. (2018) Scientific realism again. *Spontaneous Generations*. *9*, 99–107. https://doi.org/10.4245/sponge.v9i1.29356.

Ladyman, J. (2021) Structure not selection In *Contemporary Scientific Realism: The Challenge from the History of Science*, Lyons, T. D. and Vickers, P. (eds). Oxford University Press. Oxford. pp. 239–256. https://doi.org/10.1093/oso/9780190946814.003.0011.

Ladyman, J. and Ross, D. (2007) *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press. Oxford. https://doi.org/10.1093/acprof:oso/9780199276196.001.0001.

Suárez, M. (2003) Scientific representation: against similarity and isomorphism. *International Studies in the Philosophy of Science*. *17*, 225–244. https://doi.org/10.1080/0269859032000169442.

Wallace, D. (2003) Everett and structure. *Studies in History and Philosophy of Modern Physics*. *34*, 87–105. https://doi.org/10.1016/S1355-2198(02)00085-0.

Wallace, D. (2022) Stating structural realism: Mathematics-first approaches to physics and metaphysics. *Philosophical Perspectives*. *36*, 345–378. https://doi.org/10.1111/phpe.12172.

Worrall, J. (1989) Structural realism: The best of both worlds? *Dialectica*. *43*, 99–124. https://doi.org/10.1111/j.1746-8361.1989.tb00933.x.