

The Role of Reconstruction in the Elucidation of Quantum Theory*

Philip Goyal[†]

University at Albany (SUNY), NY, USA[‡]

(Dated: November 12, 2022)

I. INTRODUCTION

A physical theory must balance two very different demands. On the one hand, it must allow us to better *grasp* some aspect of the workings of the physical world; or—as it would have been common to say in a bygone era—to better understand the mind of God. On the other hand, it must actually *work*—it must provide a conceptual and mathematical framework of some generality within one can describe actual laboratory experiments and can make precise, novel predictions that conform to the brute facts of experience¹.

In the *developmental* phase of a theory, if push comes to shove, the demand for workability usually wins out. Consequently, a freshly developed physical theory is inevitably a *compromise*, which can manifest in several ways. One manifestation is that certain mathematical features of the theory’s formalism may lack clear physical motivation or meaning. Once the theory has been tested and the physics community is sufficiently convinced that its formalism captures some basic regularities in nature’s workings, there usually follows a *reflective* phase in which efforts are made to elucidate these physically-obscure features.

Reconstruction is a methodology for elucidating physically-obscure features of a theory’s mathematical formalism by deriving these features from a set of physical principles and auxiliary assumptions. An ideal reconstruction is one that traces these features back either to extant broadly-accepted fundamental physical principles or desiderata, or to newly-formulated physical principles of a widely-accepted type (such as symmetry, compositional, or extremal principles). The target of reconstruction varies according to whether one wishes to elucidate the physical basis of a specific feature of a theory’s formalism, or the formalism as a whole.

The process of reconstruction—especially wholesale reconstruction of a theory’s formalism—can be viewed as the construction of a *metatheory* that yields an existing theoretical formalism (or part thereof) as an output. From this perspective, the existing theory’s formalism is *data*, a brute fact that one seeks to understand through the principles of the metatheory. That is, reconstruction *iterates* the theory-building process: the original theory explains patterns

* To appear in ‘Phenomenology and QBism: New Approaches to Quantum Mechanics’, ed. Philipp Berghofer and Harald A. Wiltsche (Routledge Studies in the Philosophy of Mathematics and Physics)

[†]Electronic address: pgoyal@albany.edu

[‡]Homepage: <https://www.philipgoyal.org>

¹ The degree to which these demands are in fact insisted upon, or indeed are deemed appropriate, varies with time and with the physics sub-community (or indeed the physicist) in question. For example, Mach’s view of a theory as merely (or primarily) an economic codification of sense data deemphasizes the first demand, while certain modern-day research programs (such as the string theory program) appear to deemphasize the second.

in the brute facts of sensory experience; in turn, the metatheory elucidates the physical meaning of the mathematical features and structures in the thus-devised theory².

The reconstruction of a theory (or a part thereof) tends to require concepts, mathematical tools, and sometimes ways of thinking about the physical phenomena of interest, that are quite different from those that were employed in the theory's development. Accordingly, successful reconstruction of a theory must usually await the development of the appropriate concepts, mathematical tools, or new ways of thinking, and may not be achieved until many decades after the theory's formulation.

The degree to which a reconstruction enhances a theory's intelligibility depends on the extent to which the theory in question was shaped by general physical principles. In the case of Newtonian mechanics, which was substantially shaped by general principles (such as the principle of inertia and Galileo's principle of relativity), reconstructive work has tended to clarify interconnections between parts of the theory³ without shaking its deeper conceptual foundations. In contrast, Faraday–Maxwell electromagnetism was largely shaped by Faraday's imaginative and detailed engagement with electromagnetic phenomena rather than by new general principles comparable in scope to those that underpin classical mechanics. Here reconstruction had a correspondingly greater impact: Einstein's reconstruction of the Lorentz transformations—a mathematical structure that was abstracted from Maxwell's equations only decades after their formulation—led to a profound reconceptualization of the nature of space and time, and to the addition of light alongside matter in the inventory of fundamental physical entities.

Since its formulation almost a century ago, quantum theory has stubbornly resisted elucidation. It is broadly—if not universally—accepted that the theory violates numerous basic convictions about the constitution of the physical world and its relation to observers, convictions that sustained the development of classical physics for three centuries. Ideally, one would like to know what aspects of the classical conception of physical reality can be retained, what aspects need to be modified, and which abandoned; and to be in possession of an overarching conception of physical reality (analogous to the mechanico-geometric conception which underpins classical physics (Berghofer et al. 2021)) which renders these changes intelligible. However, although traditional elucidative methods (such as no-go theorems, reformulations and interpretations) have provided valuable insights, we still lack a comprehensive, compelling account of just what kind of physical reality is so extraordinarily elegantly encoded in the mathematical formalism of the theory.

In this paper, I argue that, in order to make further decisive progress, a new elucidative strategy is called for, one based on reconstruction of the quantum formalism. In particular, I propose a *two-step reconstruction-based strategy*:

1. *Reconstruct the quantum formalism.* First, reconstruct the quantum formalism, with the specific goal of distilling the full physical content of the formalism into physical principles and assumptions that can be expressed in

² This echoes Wigner's view of symmetry principles, *viz.* just as physical theories formalize regularities in sensory data, symmetry principles formalize regularities in the laws posited by those theories (Wigner 1960).

³ For example, reconstruction of the classical quantities of motion (momentum and kinetic energy) shows that these quantities are a direct consequence of Galilean relativity and the general desideratum of conservation, thereby establishing a clear connection between the *dynamical* and *kinematic* aspects of the theory, and also making clear the necessity of forms of energy not bound to massive bodies (Goyal 2020). See Sec. IV A for details.

natural language and that are amenable to philosophical reflection.

2. *Interpret the reconstruction.* Second, reflect on the principles and assumptions of the reconstruction, bringing to bear whatever philosophical traditions may be appropriate.

Ideally, the second, reflective step will yield a set physical principles and assumptions that can be laid alongside those that comprise the classical conceptual framework, enabling a point-by-point comparison which makes clear what aspects of the classical framework have been retained, modified, or abandoned. Ideally, it will also yield an overarching conception of physical reality which broadly motivates this new set of physical principles and assumptions.

This reconstruction-based interpretive strategy has many advantages over most traditional elucidative approaches. In particular, the reconstructive step potentially makes the full content of the formalism available for philosophical reflection. Hence, in the reflective step, it is possible to simultaneously take into account a larger number of the non-classical features of quantum theory. In contrast, traditional elucidative approaches take most or all of the quantum formalism as a given and typically only seek to offer explanation of specific aspects of the theory. As a result, they each harness only a small fraction of the physical content of the formalism, and generate fragmentary insights which are difficult to unify into a coherent conception of reality.

A reconstruction-based interpretative strategy is particularly timely: the *quantum reconstruction program* has galvanized the efforts of many in the quantum foundations community over the last twenty or so years, during which period several detailed reconstructions of key parts of the quantum formalism have been developed. Philosophical reflection on certain reconstructions has already been carried out, and some intriguing insights into long-standing puzzles have already recently been obtained. One of the broader aims of this paper is to stimulate the kind of collaborative work that will likely be needed to fully harvest the fruits of the quantum reconstruction program.

The paper is organized as follows. I begin in Sec. II with a discussion of classical physics, with the aim of clearly identifying what underlies the widespread view that classical physics is intelligible. I argue that the intelligibility does not consist primarily in classical physics' comportment with everyday intuitions about the physical world (as is often supposed or presumed), but rather in classical physics possessing a *coherent tripartite structure* that connects a clearly-articulated project, a conceptual framework, and specific physical theories. Accordingly, such a tripartite structure provides a template and benchmark for the elucidation of quantum theory.

In Sec. III, I discuss the lack of intelligibility of quantum theory, describing how this is a legacy of its complex and convoluted historical development. I outline the main traditional approaches for elucidating quantum theory, and analyse their limitations. I then turn to the methodology of reconstruction, and describe how it can help us to overcome the key interpretative bottleneck of the traditional approaches, thereby making the full content of the quantum formalism available to philosophical reflection.

In Sec. IV, I discuss the methodology of reconstruction *per se*, showing that it is part of the natural life-cycle of physical theories, and then summarize how it has been used in classical physics to elucidate the physical meaning of key mathematical features of both classical mechanics and electromagnetism, with the latter theory providing an excellent illustration of the potential of the reconstructive methodology for the elucidation of quantum theory.

In Sec. V, I return to the reconstruction of quantum theory, describing the operational and informational perspectives which lie behind so much of the recent success of the reconstruction program, and then summarize some recent reconstructive work.

In Sec. VI, I describe some of the key interpretational insights thus far obtained through interpretation of reconstructions of the quantum formalism, both of the abstract quantum formalism and the quantum symmetrization algorithm (which is needed for handling systems of identical particles).

I conclude in Sec. VII with some general remarks.

II. INTELLIGIBILITY OF CLASSICAL PHYSICS

The overarching goal of physics is to explore the physical world through precise, controlled experimentation; to develop a precise conception of physical reality; and on these twin bases to develop mathematically-precise, predictive theories of nature's workings. As such, the project of physics seeks to integrate three distinct facets of human activity—experimental, natural philosophic, and mathematical/theoretical. The creative interplay between these facets is most clearly visible in the development of classical physics, where they at once challenge and inspire, guide and constrain each other.

Alongside this goal is a *knowability* assumption (or article of faith) that we, no matter our limitations (in spatial and temporal reach, cognitive abilities, and experimental capacities), can—with the aid of a methodology based on the synthesis of the above three facets—aspire to a substantial knowledge of the physical world. This assumption expresses the belief or faith that reality—or a significant part thereof—is so constituted as to be amenable to discovery through a specific method.

As elaborated below, this overarching goal and knowability assumption are deeply intertwined with the fundamental assumptions that lie at the heart of classical physics and that comprise its conceptual framework. Specific theories of classical physics—such as Newtonian mechanics—are then built within that conceptual framework by proposing additional assumptions and principles appropriate to the specific group of phenomena of interest.

Thus conceived, classical physics has a tightly-linked tripartite structure:

1. *A clearly articulated project* in the form of an overarching goal and knowability assumption.
2. *A conceptual framework* which articulates a specific conception of reality, which is consonant with the overarching goal and knowability assumption.
3. *Specific theories*, with their domain-specific assumptions and mathematical predictive formalisms, which are built within the conceptual framework.

As elaborated below, I contend that this coherent tripartite structure is primary origin of our shared sense that classical physics is *intelligible*, more so than any particular way in which classical physics appears to comport with certain of our pre-scientific intuitions about the nature of reality.

A. The tripartite structure of classical physics

1. Overarching goal of physics

The above-stated overarching goal of physics combines two distinct poles—the experimental and the theoretical. The experimental ideal contains two key ideas:

E1. Isolability. One can *isolate* a physical system of interest from its past and future history, and to some extent shield it from uncontrollable contemporaneous events occurring elsewhere.

E2. Observability. An experimenter can precisely observe a system's behaviour when placed in a spatiotemporal physical context under experimental control.

Satisfaction of the isolability requirement ensures that the system of interest is insulated from that which an experimenter cannot possibly entirely know or control. Observability ensures that all aspects of the behaviour of a system are open to view, and moreover can be *precisely* measured; and that the aspects of physical context relevant to this behaviour can be controlled.

According to the theoretical ideal, an ideal physical theory has the following characteristics:

T1. Generality. The theory can be applied to all things, no matter where or when they be.

T2. Mathematical precision. The theory is cast in exact, mathematical terms.

T3. Testability. The theory generates empirically-testable predictions of sufficient precision as to permit its falsification.

The desideratum of generality focusses attention on the most deep lying regularities in events rather than the particular events themselves, and in particular seeks to eliminate any distinctions between types of physical object (for example, between the living and nonliving), their spatial locations (for example, between terrestrial and heavenly bodies) or the universal epoch during which they exist. The ideal that a theory be cast in mathematical terms and generates precise predictions establishes a tight linkage between sense and thought, while testability ensures that theoretical errors or limitations can be exposed by empirical data.

2. Classical conceptual framework

The Newtonian conceptual framework of classical physics posits that what exists is matter that moves on a stage of space in step with a universal time, with all properties of that matter open to the gaze of the ideal observer. In short, it posits a *mechano-geometric* conception of physical reality. A summary of the key assumptions underlying the Newtonian conception is given in Fig. 1, organized by those assumptions that describe (i) the space-time canvas, (ii) the matter that exists on this canvas, (iii) the motion of matter; and (iv) the observation of matter. As summarized below (Sec. II A 3), most of the assumptions can be traced back to experimental and theoretical ideals that are incorporated in the project of physics.

As classical physics developed over the course of some three hundred years, the classical conceptual framework underwent significant change—the introduction of the energetic framework, the development of the field concepts of electromagnetism, and the theories of special and general relativity, all challenged the framework’s fundamental assumptions, and it either stretched to accommodate or underwent modification. However, to keep the following discussion sufficiently focussed, these changes are not systematically considered here, although the challenge posed—and changes wrought—by electromagnetism, being of special interest in connection with the potency of the reconstructive methodology, are discussed in Sec. IV B.

i. Canvas. Assumptions C₁–C₄ license a mathematical description of space and time, so that what one experiences as an instant of time and as a location in space can be described as a region—in the limit, a *point*—in a *mathematical* manifold. Due to the assumed homogeneity and isotropy of space, the spatial portion of the manifold is ascribed particular topological and metrical structure—it is assumed to be simply-connected and Euclidean—and is assumed to be three dimensional. Similarly, due to the homogeneity of time, the temporal part of the manifold is ascribed a uniform metric. Thus, the underlying canvas is highly symmetric, bereft of any structure corresponding, for example, to the spatial and temporal asymmetries—*up* vs. *down* or *past* vs. *present* vs. *future*—so characteristic of lived experience, although spatiotemporal continuity and three-dimensionality, for instance, seem to be reasonable extrapolations of our lived experience.

ii. Composition of bodies. Assumptions P₁–P₅ reflect the atomistic conception of physical bodies, namely that every body can be reduced to an arrangement of particles. Bodies are assumed to be *composed* of these particles in a manner that precludes strong emergence (i.e. the ‘whole’ is nothing more than the arrangement of its ‘parts’) (P₁). The particles themselves are assumed to be indefinitely persistent (they are not created or destroyed) (P₂), which can be viewed as arising from the desideratum of maximizing symmetry (the existence of particles of finite lifetime would introduce an inherent distinction between different moments in time). Moreover, all particles are assumed to be characterised by the values of just a few property-types, such as mass and charge (P₃–P₅), with particle position being a particle’s only *directly* observable property (O₃) (all its other properties are dispositional, manifest to an observer only indirectly through their effect on position in certain experimental contexts).

iii. Motion. Assumptions M₁–M₆ concern the motion of matter, and range from the general to the specific. The twin requirements of determinism and reversibility (M₁), together with continuity (M₂), ensure that the persistent particles which compose all bodies are *reidentifiable*—that is, each particle has its own trajectory throughout time, along which it can (due to O₁, O₂, and O₃) be precisely tracked. If reversibility (M_{1b}) were not to hold, it would be possible for two particles’ trajectories to merge, which would prevent their reidentification were they to be identical in their time-independent properties. Similarly, were continuity (M₂) not to hold, so that particles could ‘jump’ from place to place, then reidentification would be rendered at best probabilistic.

The assumption that an isolated particle moves uniformly in a straight line (M₃) establishes a direct connection between the posited mathematical structure of space-time and the behaviour of particles immersed therein—isolated particles move equal spatial distances in equal times (as judged by the spatial and temporal metrics previously posited).

The formulation of the specific equations of motion of classical mechanics depends centrally on continuity (M₂),

conservation (M₄), and differential composability (M₅). Conservation posits that certain physical quantities remain unchanged during such fundamental processes as the collision of two particles, and provides the basis for the later introduction of specific *quantities of motion* (such as kinetic energy and momentum). Continuity enables one to specify motion via *differential* equations, which specify how a system will change in an infinitesimal of time. And differential composability posits that, in an infinitesimal of time, the change in motion of a given particle in a system of particles can be understood as a composition of the changes due to the two-particle interactions between that particle and each of the other particles in the system. This assumption enables the ‘scaling up’ of the equations of motion for a two-particle system to a system of any number of particles (in Newtonian mechanics, the assumption leads to a rule for the composition of forces) (see Sec. II A 4).

iv. Observation. Assumptions O₁–O₄ connect the picture of physical reality painted by the previous sets of assumptions (canvas, composition, motion) to that which can be observed. Throughout, the notion of *observation* is taken as primitive. Operationally, observation can be treated as an abstraction of the experiences of an idealized experimenter. Ontologically, it could be viewed in several ways, for example (i) as reflecting the action of an abstract mind or consciousness of unspecified nature capable of becoming *aware* of (and thus *observing*) outcomes of measurement devices, or (ii) as reflecting the actualizing of events (outcomes of measurement devices) that are themselves taken as primitive (and as distinct from material objects and their properties).

Assumption O₁ posits that the ideal process of observation is *passive* (or sight-like) in the sense that an ideal measurement of a particle’s properties *registers* those properties without bringing about any change in their values. This assumption ensures that one can regard a given *appearance* as a direct manifestation of a pre-existing property—for example, the appearance of a particle at a location in space as underpinned by the body *possessing* (immediately prior to the observation) the property of being located at that location in space.

Assumption O₂ posits that particle properties, which are quantifiable (P₅), are *exactly* observable. Assumption O₃ allows for the existence of idealized rods and (together with M₃) idealized clocks, thereby furnishing observational access to the posited space and time structure insofar as spatial *intervals* and temporal *durations* are concerned. An idealized clock is one that indefinitely beats at a constant rate (irrespective of its location or state of motion), while an idealized rod is one that indefinitely possesses one and the same length (irrespective of its position, orientation, or state of motion). Note that the homogeneity (C₂) and isotropy (C₃) of space, together with inertial-frame equivalence (O₄), precludes measurement of *absolute* location, direction, and velocity, while temporal homogeneity (C₄) precludes measurement of absolute time. Assumptions O₂ and O₃ jointly imply that the measurement of all particle properties can be accomplished via suitable position measurements.

Finally, assumption O₄ is a specific assumption, connected to M₃, which asserts that all inertial frames (namely those in which an isolated body is observed to travel rectilinearly at constant speed) are ‘physically equivalent’ in the sense that the laws of motion apply equally to any such frame.

Observation vs. Manipulation. The notion of passive observation (O₁) is in tension with an experimenter’s capability—which he necessarily possesses—of *manipulating* physical bodies (such as experimental equipment) *at will*. Here, ‘will’ is taken as a primitive and is modelled within the mechanical conception of causation as that which can *initiate* causal

chains.

Given that the experimenter is able to manipulate matter at will, questions concerning the nature of this will cannot be entirely separated from a description of the nature of the physical world. In particular, what constraints apply to ‘at will’ manipulation? Does the will act only through the body of an experimenter? Do other kinds of physical systems manifest a will? Is the will of the experimenter influenced by the physical system of interest (for example by its state)? If the will is not wholly independent of the physical state of the system of interest (or the physical world *in toto*), as one would expect from experience, then the existence of the experimenter in the physical world generates a complex feedback between the otherwise-deterministic physical world and the experimenter, blurring the line between the physical and ‘non-physical’.

The classical conceptual framework provides no fundamental answers to these questions, in particular providing no explanatory grounds for our experience as causally-*efficacious actors* in the physical world. Nevertheless, the framework posits that a physical system subject to observation *in the absence of manipulation* evolves deterministically and reversibly. Hence the Newtonian metaphor of the solar system as a clockwork mechanism holds good as long as one excludes such agential manipulations as the deflection of asteroids by explosive detonations.

3. *Relation between overarching goal of physics and the classical conceptual framework*

The experimental and theoretical ideals described above (Sec. II A 1), in combination with the assumption of knowability, are supported by the conceptual framework of classical physics (Sec. II A 2). Below, the experimental and theoretical ideals, and the fundamental assumptions that underpin the conceptual framework, are referenced by their labels (C1, etc.).

1. The ideal of *isolability* (E1) is supported by:

- (i) a particular view of causation, namely that the behaviour of a system may be understood via efficient cause, without recourse to final cause. Accordingly, the future trajectory (or any aspect thereof) of the system is irrelevant to the present, provided that sufficient cognizance is taken of its past. Moreover, only its *immediate* past is relevant—its more distant history is rendered irrelevant by sufficient knowledge of its immediate past (M1a).
- (ii) the assumption that a system is disproportionately influenced by bodies in close physical proximity (M6), or that influences due to more distant bodies can be regarded as providing a stable background for experimentation.

2. The ideal of *observability* (E2) is supported by the assumption that the idealized process of observation has negligible impact on the system under observation, and moreover that the system and its properties *exist* and are well-defined quite independently of the process of observation. That is, idealized observation is *sight-like* rather than *touch-like* (O1).

Canvas

- C1.** *Canvas quantifiability.* Space and time are quantifiable.
- C2.** *Spatial homogeneity.* Space is homogeneous.
- C3.** *Spatial isotropy.* Space is isotropic.
- C4.** *Temporal homogeneity.* Time is homogeneous.

Composition of bodies

- P1.** *Simple mereology.* Bodies are simply-composed of particles.
- P2.** *Particle eternity.* Particles are indefinitely persistent.
- P3.** *Property universality.* Particle property-types are universal.
- P4.** *Property sparsity.* Particle property-types are few.
- P5.** *Property quantifiability.* Particle properties are quantifiable.

Motion

- M1.** *Determinism and Reversibility.* Motion is (a) deterministic and (b) reversible.
- M2.** *Continuity.* Motion is continuous.
- M3.** *Isolated motion.* An isolated body moves uniformly and rectilinearly.
- M4.** *Conservation.* Collective motion conserves certain quantities of motion.
- M5.** *Differential composability.* Collective motion is differentially composable.
- M6.** *Influence fall-off.* Interparticle influences monotonically diminish in strength with interparticle separation.

Observation

- O1.** *Sight-like observation.* Ideal observation passively registers particle properties.
- O2.** *Exact observability.* All properties of bodies are exactly observable.
- O3.** *Position categoricity.* Particle position is the only directly accessible particle property.
- O4.** *Principle of relativity.* All inertial frames are physically equivalent.

FIG. 1: *Conceptual framework of classical physics.* A list of the fundamental assumptions underlying classical physics in the Newtonian era, up to (but not including) the development of electromagnetism.

3. Due to *generality* (Γ_1), one aims at a theory that describes a ‘bottom’ layer of physical reality that is *maximally symmetric*, and in which there are a minimum of primitive entities. In particular, one aims to:
 - (i) treat all places, times, frames of observation, and objects—as far as possible—in the *same way*. For example, space is viewed as homogenous and isotropic (C_2 , C_3); time as homogeneous (C_4); all inertial frames as equivalent (O_4); and all objects—whatever their visible form—as describable in terms of the same small number of basic properties (P_3 , P_4) whose values can gleaned through position measurements (O_3).

- (ii) treat ‘nonfundamental’ things as simple composites of ‘fundamental’ things. For example, regard macroscopic bodies as simple composites of *particles* (P₁) which can be drawn from a small number of possible types (an expression of Greek atomism). Here, simple composition precludes strong emergence.
4. The ideals that a theory be expressed *mathematically* (T₂) and that the behaviour of a system can be *precisely* observed (E₂) are supported by the assumption that particle properties are quantifiable (P₅) and can be exactly observed (O₂). Since our visual sense is the most readily and unambiguously quantifiable, that in turn favours the view that all properties can be mapped (using suitable experimental and theoretical tools) to spatial extension (O₃), and that spatial intervals and temporal durations are observable (O₃, M₃). Mathematical expressibility also favours:
- (i) *domain-specific principles*: uniformity assumptions such as the indefinite persistence (eternality) of the fundamental constituents of matter (P₂) and dynamical uniformity principles (such as uniform motion of isolated bodies (M₃) and conservation principles (M₄)), as well as other dynamical principles (for instance extremization principles like the principle of least action) that can readily be expressed in mathematical terms; and
- (ii) *technical assumptions*: such as continuity of trajectories (M₂), and the existence of symmetries (such as associativity and commutativity of composition).
5. The ideal of *testability* is supported by the assumptions of determinism and reversibility (M₁), since both promise the most perfect conceivable level of prediction. This ideal also disfavors theories that posit quantities (such as ‘hidden variables’) or detailed structure which is not directly amenable to experimental test.

4. *Theories of classical physics*

Newtonian mechanics, the first comprehensive theory of classical physics, is formulated within the foregoing classical conceptual framework. In brief, the theory introduces a crucial distinction between time-independent and time-dependent properties, positing that particles have a time-independent mass, together with two time-dependent properties—position and velocity. On this basis, conservation (M₄) can be mathematically expressed as the conservation of total momentum in collective motion. The interactions of two separated bodies is handled through Newton’s concept of *force* together with a specific law connecting force and change in velocity, with the posited inverse-square gravitational force between two bodies satisfying the influence fall-off assumption (M₆). Differential composability (M₅) is then formalized via the parallelogram rule for the composition of forces, enabling the model of two-body systems to be ‘scaled up’ to systems of any number of bodies⁴.

⁴ In Goyal (2020), Newtonian mechanics is reconstructed on the basis of these and other special assumptions. A categorization of the assumptions employed is also given.

Other fundamental theories of classical physics—chiefly electromagnetism and Einstein’s relativity theories—proceed analogously, introducing additional domain-specific assumptions that posit specific quantities of interest and that enable the construction of a set of differential equations of motion describing how these quantities change over time. As mentioned above, each such theory has challenged the Newtonian-era classical conception, which has either stretched to accommodate or has been modified appropriately. Nevertheless, broadly speaking, the fundamental theories of classical physics—Newtonian mechanics, Faraday–Maxwell electrodynamics, and Einstein’s relativity theories—are all recognizably *classical* insofar as they share most of the same basic convictions as to the nature of the physical world and the nature of observation.

B. Why is classical physics intelligible?

It is frequently suggested or supposed that the intelligibility of classical physics is a rather trivial consequence of the fact that it reflects ‘common sense’—that its conceptual framework is intelligible primarily *because* is a refinement and abstraction of our mental model of an everyday physical world, namely a world populated by persistent objects immersed in space which possess various properties (like location, shape, colour), move smoothly from place to place, and which we observe (chiefly through the sense of sight) largely as they really are.

While it is certainly true that classical physics inherits many of these commonsense notions, it is also the case that, in many respects, it departs from everyday notions and intuitions. For example, whereas the canvas assumptions posit an isotropic space (C₃), our lived experience reflects an asymmetry of ‘up’ and ‘down’. Similarly, the posit of homogeneous time (C₄) together with determinism and reversibility of motion (M₁) erase any theoretical distinction between past, present and future, and makes it difficult to understand, in *fundamental* terms (i) why we do not experience the entire history of the universe in one go but only the *present moment*; (ii) what is the nature of genuine creativity (or genuine novelty) or a sense of purpose that is future-directed. And the assumption of simple mereology (P₁) reduces all objects to mere arrangements of particles, depriving a ‘whole’ with any fundamental meaning other than the arrangement of its parts, for example erasing any *fundamental* distinction between the living and nonliving.

Ideally, the project of physics would wish to *contingently recover* these intuitions as far as possible—to explain just how it can be the case that we have these intuitions in spite of the validity of the classical conceptual framework by referring to specific contingent facts. In a few cases, this can be said to have been achieved. For example, it is now widely accepted that our sense of the asymmetry of up and down is a consequence of our contingent Earth-bound existence, and that this asymmetry would effectively disappear if we were to live ‘on the float’ on board a non-rotating space station far from gravitating bodies. However, in most other cases—such as in the distinction between past, present and future; between the living and nonliving—no such robust reconciliation can be said to have been achieved, notwithstanding the enormous insights that have in some instances been gained by proceeding in the hope or belief that such a reconciliation might be possible.

For these reasons, I assert that the intelligibility of classical physics primarily reflects the fact that it abstracts

some particular aspects of human experience in such a way as to create a tripartite structure which tightly weaves a clearly-articulated project, a well-defined conceptual framework, and empirically-successful mathematical theories of specific physical phenomena. From this perspective, classical physics is the product of what may well be a necessary *trade-off*: between the desideratum of creating general, mathematically-precise, predictively-capable physical theories, on the one hand; and the desideratum of taking full account of the vast richness of the human experience of that which we call reality, on the other. In this sense, classical physics is a bright spotlight that brightly illuminates—and brings into sharp relief—some aspects of our experience of reality, whilst leaving the substantial remainder in shadow.

III. INTELLIGIBILITY OF QUANTUM THEORY

A. Development of quantum theory

What is commonly referred to as *quantum theory*—the von Neumann–Dirac abstract quantum formalism at its core—is the end-point of a period of development that spanned some twenty-five years. From the first to the last, this development was marked by bold conjectures that were made, above all, in the quest to construct viable models of specific groups of microphysical phenomena that, until that point, had entirely eluded models which squarely obeyed the strictures of classical physics. Almost without exception, these conjectures lacked any compelling *a priori* physical justification or oftentimes any justification at all. Sometimes they were directly at odds with classical physics. But, by and large, they were taken seriously because they *worked*.

For example, at the start of this period, in 1900, Planck noticed that if one posits that the energy of certain abstract oscillators is *quantized*, one can use then-standard statistical arguments to derive the blackbody radiation curve. This assumption was essentially an *ad hoc* mathematical device, whose sole justification was that it produced a result consistent with known experimental data, something which classical models had completely failed to do. No *a priori* physical rationale for such quantization was offered, nor was one on the horizon—classical physics embraced continuity. Similarly, in order to make theoretical contact with the discrete pattern of spectral lines exhibited by hydrogen, Bohr freely combined Newtonian mechanics and electrostatics with a novel assumption—the quantization of angular momentum—inspired by Planck’s, which gave rise to a discrete set of electronic orbits. In this case, the novel assumption was not only a freestanding assumption (*viz.* apparently not derivable from classical physics and not provided with a deeper rationale) but was at odds with an established theory of classical physics—electromagnetism predicts that an accelerating charge radiates away its energy, so on this account no stable electronic orbits should exist.

This pattern continued up to and including the formulation of a general mathematical formalism of quantum theory in the mid-1920s. For example, Schroedinger’s wave equation for a single particle (such as an electron), itself based on de Broglie’s bold conjecture of wave–particle duality, was originally conceived as analogous to a classical wave equation. But, although Schroedinger initially proposed to consider the wave as describing a spatial charge density, it soon became clear that, in order to make proper connection between the wave equation and the brute fact of particle-like detections in the laboratory, one had to supplement the wave equation with a separate rule—the Born

rule—that gives the *probability* that a particle will be detected within a region of space. Similarly, in order to handle systems of identical particles, such as multielectronic atoms, Dirac and Heisenberg independently introduced novel mathematical rules (the symmetrization postulate, in particular) which they rationalized by claiming that identical particles are *indistinguishable* from one another. Both of these ideas were at odds with basic tenets of the classical conceptual framework—probabilistic measurement outcomes in conflict with the sight-like nature of classical measurement and with the knowability assumption; indistinguishability in conflict with classical reidentifiability.

By the end of this period of development, the broad consensus in the physics community, particularly amongst most of the founders of quantum theory, seems to have been that the new theory departed from classical physics to such an extent that it was neither possible nor fruitful to try to stretch or modify the classical conceptual framework to accommodate it. Rather, it was necessary to somehow develop a new conceptual scheme that would render the mathematical formalism intelligible. However, given the huge cognitive cost of setting aside the classical framework—a framework that was widely regarded as *intelligible* and moreover had sustained the development of physics for more than two hundred and fifty years—there was understandably considerable dissent, including from such key figures as Einstein and Schroedinger.

Despite the substantial, multifaceted efforts made to elucidate quantum theory in the intervening period of roughly one hundred years, quantum theory is still broadly regarded as mysterious and counterintuitive, amongst both physicists and other academicians engaged in its use or conceptual investigation (such as other natural scientists, philosophers of physics, and philosophers of science), as well as amongst those members of the wider academy and general public interested in the broader implications of natural science.

B. Rendering quantum theory intelligible

It is commonly asserted (or supposed) that quantum theory is unintelligible because it violates so many of our everyday intuitions, particularly those that seem to be conditioned by our engagement with the macroscopic physical world. There is undoubtedly some truth in this statement. But, from the perspective of the analysis of the intelligibility of classical physics given above, the unintelligibility of quantum physics does not *primarily* originate in its departure from ‘common sense’. Rather, the lack of intelligibility is a symptom of quantum theory not being embedded in a conceptual framework—a *quantum conceptual framework*, if you will—that makes explicit *what* aspects of human experience it abstracts from and quite *how* it abstracts from them.

However, it is evident that the efforts made thus far to elucidate quantum theory have somehow fallen short—notwithstanding the valuable insights they have undoubtedly provided, they have not sufficed for the construction of a comprehensive quantum conception of physical reality. I contend that, in order to make decisive further progress, we need a better *methodology* for elucidating quantum theory, and that a suitable methodology is one that leverages the *quantum reconstruction program*, a program that has attracted considerable attention in the quantum foundations community over the last twenty or so years.

Below, we first briefly review the main approaches that have traditionally been used to elucidate quantum theory,

and identify the nature of their limitations. We then turn to the methodology of reconstruction, and examine its potential for helping us develop a thoroughgoing understanding of quantum theory.

C. Traditional approaches to elucidate quantum theory

Broadly speaking, the most prominent and influential approaches traditionally employed to elucidate quantum theory are as follows:

1. *New concepts*. Formulate new concepts that encapsulate fundamental ways in which quantum theory seems to be at odds with the classical conception of physical reality, and then investigate quantum theory quantitatively through these new conceptual lenses. Examples include complementarity and entanglement.
2. *No-go theorems*. Show that certain predictions of quantum theory are inconsistent with a minimal classical model. Notable examples include Bell's theorem (nonlocality) (Bell 1964) and the Kochen-Specker theorem (noncontextuality) (Kochen and Specker 1967).
3. *Reformulations*. Re-write some or all of the quantum formalism in an alternative mathematical form. This often establishes an illuminating parallel between (parts of) quantum theory and an existing mathematical or physical theory. Examples include Feynman's path-based reformulation (which establishes a parallel to probability theory) (Feynman 1948), and Bohm's Hamilton-Jacobi formulation of the Schroedinger equation (which establishes a parallel to the Hamilton-Jacobi formalism of classical mechanics) (Bohm 1952).
4. *Interpretations*. Formulate a conceptual framework which at least partially accounts for some of the nonclassical features of the quantum formalism. Examples include Bohr's complementarity-based interpretation (Bohr 1928), Heisenberg's potentiality/actuality-based interpretation (Heelan 2016), the Stapp-Schwarz mind-based interpretation (Schwarz et al. 2004), the transactional interpretation (Cramer 1986), the many worlds interpretation (Everett 1957), the de Broglie-Bohm interpretation (Bohm 1952), and the QBist interpretation (Fuchs 2017).

A brief remark: It is important to bear in mind a distinction between a *research program* and a *component* of it. For example, understood as a research program, Bohmian mechanics contains both a reformulative and interpretative component (both of which components are mentioned above)⁵. But it also involves other components, such as ongoing attempts to derive the quantum equations of motion from more basic assumptions, or attempts to test the interpretation⁶.

⁵ To be clear: the reformulative component consists in a re-expression of parts of the quantum formalism, for example the reformulation of the Schroedinger equation in the form of a Hamilton-Jacobi equation. Minimally, the interpretative component consists in the the posit of the existence of particles obeying a guidance condition, together with an interpretation of the wave function. These two components are tightly connected, but can be separated—the reformulation is a purely mathematical rewriting, without any additional ontological commitment. And the reformulative component is being actively developed, for example being extended to relativistic quantum mechanics and quantum field theories.

⁶ For example, through establishing whether a $|\psi^2|$ -distribution of particles can be viewed as an 'equilibrium' distribution by showing that a 'non-equilibrium' distribution—i.e. one not conforming to $|\psi^2|$ —will 'relax' to an equilibrium distribution under unitary dynamics.

Similarly, the many worlds research program has not only an interpretative component, but also a reconstructive component (which one could view as part of its attempt to *test* or bolster the interpretation) consisting of a derivation of the Born rule from decision-theoretic axioms. And the QBist research program has—in addition to its interpretative component—reformulative (re-expressing the quantum formalism as far as possible in terms of SICs and associated probabilities) and reconstructive components.

However, it is usually the case that a research program is ‘known’ mostly for one of its components. Some of these components may be relatively stable, while others may be in active development. Research programs also differ in the degree of cohesion between their components. These nuances would need to be taken into account by an in-depth discussion of any of the research program components mentioned above. Here, I am focussed on the relative merits and limitations of the methodologies that have been traditionally employed for elucidating quantum theory, rather than on entire research programs *per se*. So, the focus is on components. I shall mention the connections to other components in a research program only if directly relevant.

1. Illuminative value of traditional approaches

Each of these approaches has proven itself capable of illuminating quantum theory in its own distinct way:

1. The coining of new concepts is a crucial first step to elucidation, insofar as such concepts *point* to distinct non-classical features of quantum theory that appear to be of fundamental importance. Bohr’s concept of *complementarity* sought to provide de Broglie’s bold conjecture of wave-particle duality with a philosophical underpinning, arguing that it reflected an irreducible trade-off between space-time coordination (specifying *when* and *where* an object is located) and causality (the possibility of connecting the future behaviour of an object with its past behaviour). The concept of complementarity has since been formalized in many different ways, such as in the context of POVMs, information trade-off relations, and mutually unbiased bases.
2. No-go theorems, such as Bell’s theorem, establish that, in light of quantum theory, we must abandon one or more classical assumptions about the nature of physical reality or our access to it. As such, these theorems force us to recognise that certain patterns of thinking are not viable, and spur efforts to find theoretical or experimental loopholes, and to develop alternative patterns of thinking that are in tune with quantum theory.
3. The abstract quantum formalism is standardly expressed in the language of complex vector spaces, with physical states represented as complex vectors (or rays), dynamics as unitary operators, and projective measurements as Hermitian operators. Re-expression of this formalism, or of specific quantum models, in mathematical garb that establishes parallels with an existing mathematical or physical theory can allow us to appreciate quantum theory from another angle, and often suggests new ways of understanding certain nonclassical features, as well as sometimes suggesting alternative computational techniques. For example, Feynman re-expressed the abstract quantum formalism in a manner that parallels the Lagrangian formulation of classical mechanics, which results in dramatic mathematical simplification, and establishes a close parallel to the formalism of probability

theory. Similarly, Bohm expressed Schroedinger's equation in the form of a classical Hamilton-Jacobi equation supplemented by a so-called quantum potential. The Bohmian model for a system of two particles renders nonlocal influence explicit via the quantum potential, which directly inspired Bell's theorem.

4. Each interpretation of quantum theory tends to focus on making sense of some specific nonclassical features of the quantum formalism. Most are concerned with making sense of the quantum formalism's distinction between dynamics and measurement, or with the limited access that measurement provides to the degrees of freedom in a quantum state, but some (such as the transactional interpretation) are more focussed on the puzzle of nonlocality. Some interpretations seek to 'defuse' these nonclassical features, somehow stretching the classical framework to accommodate them; while others seek to provide some kind of deeper understanding of these features. In so doing, these interpretations bring attention to specific noteworthy aspects of the formalism, and provide a certain way of thinking about them, which can be creatively inspiring⁷.

2. *Limitations of traditional approaches*

However, these traditional approaches are limited in fairly obvious ways.

i. Limitations of no-go theorems. Although no-go theorems are powerful in that they force us to recognize that certain sets of obvious-seeming assumptions are logically at odds with certain quantum theoretic predictions, they are essentially *negative* results—they block certain patterns of thinking, but provide little positive guidance on how we *could* think about physical reality so as make sense of these predictions. In addition, as these no-go theorems make use of *specific* quantum theoretic predictions rather than the theory *as a whole*, the insight that they provide is necessarily fragmentary—each theorem focusses on some specific nonclassical feature—so one is left with the daunting challenge of integrating these fragmentary insights.

ii. Limitations of traditional interpretations. Interpretations of quantum theory are, in contrast, relatively *positive* in the sense that they offer a specific way to make sense of some nonclassical features of the theory. However, they only harness a small fraction of the physical content of the quantum formalism: they focus only on making sense of some specific part of the quantum formalism—typically the Born rule—and offer no explanation of the remainder, which they take as a given.

For example, the many worlds interpretation takes the unitary part of the abstract quantum formalism (that is, the abstract quantum formalism apart from the Born rule) as a given. As a result, the interpretation cannot, due to its starting point, offer an explanation of the unitary part of the quantum formalism—for example, why quantum states are represented by complex vectors and dynamics by unitary operators. Conversely, the interpretation is not

⁷ For example, David Deutsch, one of the founders of the field of quantum computation, credits the many worlds interpretation as providing a way of thinking which inspired the idea that it might be possible to *use* 'parallel universes' to carry out certain computational tasks more effectively than a classical computer. Similarly, John Bell has cited the de Broglie-Bohm interpretation as an inspiration for his eponymous theorem. These examples illustrate the powerful heuristic value of attempts—however inadequate or incomplete—to penetrate the quantum veil.

constrained by these mathematical structures—as far as the interpretation is concerned, quantum states and dynamics could be differently represented. Similarly, the de Broglie–Bohm interpretation takes the Schrodinger equation as a given, which means that its so-called quantum potential (which effectively encodes all nonclassical behaviour) must be taken as a given—as a brute fact—within the confines of this interpretation.

A further serious difficulty with traditional interpretations is that they are difficult to subject to scientific test, whether theoretical or experimental. The de Broglie–Bohm interpretation fares better than most in that it is based on a *reformulation* of quantum theory which casts parts of the quantum formalism in a mathematical form that more closely parallels classical mechanics. Consequently, one way to build confidence in this type of reformulation is to successfully extend it to other parts of the formalism, such as relativistic quantum theory or quantum field theory. However, in other cases, the possibility for such theoretical test is much more limited. For example, the many worlds interpretation offers only one obvious test, namely the derivation of the Born rule⁸. But, again, as the unitary part of the abstract quantum formalism is taken as a given, one cannot test the interpretation by asking for a derivation of that. More seriously, *experimental* tests of current interpretations seem almost entirely out of reach.

In summary, the plethora of interpretations—which cover a wide range from those that hew closely to the classical conceptual framework, to those which draw upon metaphysical or psychological notions—show that it is possible to clothe quantum theory in many different conceptual outfits in a manner that might be attractive (or even creatively inspiring) to some. However, it has proven difficult to devise compelling tests (theoretical or experimental) of these interpretations. Consequently, one has little evidential basis to choose between them. Moreover, as traditional interpretations take most of the quantum formalism as a given, they make very limited use of the actual physical content of that formalism.

D. A new approach: Reconstruction-based interpretation of quantum theory

1. The bottleneck: taking the quantum formalism as the starting point

As noted above, almost all traditional methods for elucidating quantum theory *start* with the mathematical formalism of quantum theory. They then reflect upon it, derive critical predictions from it, reformulate it, or interpret it. But, as summarized in Sec. III A, the quantum formalism was the end-point of a rather convoluted and complex process involving many novel physical ideas and a good deal of mathematical guesswork. As a result, quite what any part of the mathematical formalism really *means* is often unclear. Moreover, even in those cases where the meaning seems ‘obvious’, it is perilously easy to be misled.

For example, it is common to talk about *quantum states* and *particles*, but it is far from clear that quantum states are anything like the states of classical systems, in the sense of being a description of their *objective* physical states. And it

⁸ A decision-theoretic argument for the Born rule has been offered, but as the many worlds interpretation does away with the idea that an agent has one and only one future successor, the applicability of decision-theoretic axioms such as a scenario is questionable. In particular, in a branching universe, the notion of probability has to be reconceived.

is far from clear whether so-called quantum particles have *any* of the fundamental characteristics (such as continuous localization, persistence, and reidentifiability) that we attribute to classical Newtonian particles. Yet, it is all too easy for one to be unconsciously influenced by the classical associations conjured up by such language, in part because it is yet to be established which associations *do* carry over and which do not.

There are also more technical examples of considerable importance where one is easily misled by classical associations. For example, in the interpretation of the quantum rules for handling a system of identical particles, a great deal turns on how one reads the indices in the (anti-)symmetrized states that describe such a system. These indices have almost universally been read as *particle labels*, which is a carry-over from classical physics. Yet it has recently become increasingly apparent that this ‘common sense’ reading (recently dubbed ‘factorism’ (Caulton 2014)) may well be incorrect. But, in that case, quite what these indices *do* mean remains controversial, yet a great deal (such as the nature of entanglement in systems of identical particles) turns on this issue.

To be clear: at the instrumental level, the formalism is sufficiently well defined that it enables physicists to build models of *most* specific quantum systems of interest⁹. It is this capacity, and the empirical success of the resultant models, which justly confers such prestige upon quantum theory. However, extracting any reliable information from that formalism as to the nature of the underlying physical reality is fraught with difficulty. As mentioned above, the no-go theorems are the most reliable way of doing so, but these necessarily suffer from fragmentariness due to the fact that they only make use of specific quantum theoretic predictions rather than the quantum formalism as a whole.

2. *A remedy: Reconstruction followed by interpretation*

The methodology of quantum reconstruction seeks to remove the interpretative bottleneck by systematically deriving the quantum formalism in an operational framework from postulates that are, ideally, physically well-motivated, thereby *distilling* the full mathematical content of the theory into precise natural-language statements that—unlike the abstract mathematical postulates of quantum theory—are amenable to philosophical reflection (Berghofer et al. 2021, Fuchs 2002, Grinbaum 2007b).

Such a reconstruction can then serve as a stepping stone to interpretation: with a reconstruction in hand, one can *interpret quantum theory by reflecting on the postulates of the reconstruction, rather than attempting to decipher the inscrutable abstract mathematical postulates in which quantum theory is standardly cast*. In contrast to traditional interpretations, such an interpretation would likely be more *reliable* since the postulates of the reconstruction would likely be couched in language that is closer to basic experimental operations and be expressed in simpler mathematical terms. And it would likely be more *comprehensive* in that, owing to the greater digestibility of these postulates, it would likely be easier to simultaneously take into account a larger number of the nonclassical features of quantum theory.

In summary, then, we propose a two-step approach to interpretation:

⁹ But, even at this level, it should be noted that the standard quantum formalism is sometimes ambiguous. For example, given the standard interpretation of the quantum symmetrization algorithm for handling systems of identical particles, it is unclear how to determine whether or not a given state is entangled (Ghirardi et al. 2002).

1. *Reconstruct the quantum formalism.* Systematically derive the mathematical formalism of quantum theory in an operational framework from postulates whose physical meaning is as clear as possible.
2. *Interpret the reconstruction.* Philosophically reflect upon the postulates of the reconstruction, bringing to bear whatever philosophical traditions or notions seems fruitful.

A few remarks are in order:

1. *Operational framework.* As illustrated above, traditional interpretation of the quantum formalism suffers from considerable ambiguity in that it is far from clear to what extent concepts (such as ‘state’, ‘particle’) and mathematical features (such as indices in symmetrized states) have the meaning which they are usually ascribed in classical physics. One way to ameliorate such difficulties is to reconstruct the quantum formalism within an *operational framework*—an idealized representation of the experiments carried out on a laboratory workbench. As elaborated in Sec. VB, concepts such as physical system, measurement, and interaction, are taken as primitive, and understood in terms of concrete macroscopic devices and their observable outcomes. Proceeding in this manner, the additional concepts that one introduces in the process of reconstruction can be directly related to elementary experimental operations, which strongly constrains their meaning.

2. *Intuitively graspable postulates.* For a reconstruction to be *suitable* for interpretation, it is essential that its postulates be formulated with a view to interpretation—as far as that is possible given the severe challenge of devising a viable reconstruction (namely one that leads to some part of the quantum formalism). This speaks against postulates that are rather abstract or mathematical, and speaks in favour of postulates that express intuitively graspable ideas. In this regard, the principles that underlie classical mechanics—such as the principles of relativity and conservation, and Newton’s action-and-reaction principle—are exemplars, with the caveat that the postulates needed to reconstruct quantum theory will undoubtedly be quite different in character and may severely challenge our customary patterns of thought.

3. *Reconstruction of full quantum formalism.* What is ordinarily referred to as ‘quantum theory’ in fact consists of several distinct components (see Fig. 2), all of which are necessary to create explicit quantum models of physical systems of interest. Since it is only these explicit models that have been subjected to experimental test, it is essential for the construction of a *full* picture of quantum reality that *all* of these components be reconstructed and subjected to interpretation.

4. *Form of ideal interpretation.* What should an ideal ‘interpretation’ of a quantum reconstruction look like? I contend that, to ensure maximal intelligibility (along the lines of classical physics, as discussed in Sec. II B), such an interpretation must provide:

- (i) *A quantum conceptual framework.* This framework should be articulated in the form of a set of fundamental assumptions which one can put alongside those that underpin the classical conceptual framework. This will enable a point-by-point comparison of the two conceptual frameworks.

Abstract quantum formalism (aQF)

- **Single systems.** Mathematical representations of the physical states and temporal evolution of a physical system, and of measurements performed upon it. Representation of symmetry transformations of the frame of reference.
- **Composite systems.** Tensor product rule for a system composed of nonidentical subsystems in the case where the subsystems are in pure states.

Quantum symmetrization algorithm (QSA)

- Rules for constructing the states (the symmetrization postulate) and measurement operators for a system composed of identical subsystems (i.e. subsystems with the same time-independent properties).

Spin-statistics connection (SSC)

- Rule specifying the connection between an identical particle's spin and the applicable 'statistics'.

Quantization rules (QRs)

- **Heisenberg's equation.** General form of time-evolution unitary operator in terms of the quantum Hamiltonian operator (*viz.* the operator that corresponds to what we classically understand as a "measurement of energy").
- **Operators for specific measurements.** Operators representing measurements classically described as measurements of function of classical observables, given the operators corresponding to the latter observables. Canonical commutation relations. Explicit representation of fundamental measurement operators.

Quantum wave equations (QWEs)

- Single-particle wave equations, especially the Schroedinger, Dirac and Klein-Gordon equations.

FIG. 2: *Main components of the standard quantum formalism.* The abstract quantum formalism (aQF) is an abstract mathematical *shell* in which quantum models of specific systems of interest can be built. To construct a model of the electron in a hydrogen atom, one must employ Heisenberg's equation, and then appeal to the operator rules to establish the explicit form of the operators that represent measurements of position and momentum (this requires that one assume the classical Hamiltonian for a charged particle in an electromagnetic field). For a system composed of more than one identical particle—such as the helium atom or the conduction electrons in a metal—one also requires the quantum symmetrization algorithm (QSA) and the spin-statistics connection (SSC). Finally, to adequately describe a single particle in the relativistic regime, one requires specific wave equations, most importantly the Dirac equation for the electron. Excluded here, for brevity, are the rules associated with quantum field theory (QFT).

- (ii) *A revised notion of the project of physics*, in which the new conception is firmly rooted. For example, if quantum measurement is taken to be touch-like, so that observation is regarded fundamentally as an active process, then it is essential that this assumption be rationalized and rendered intelligible through a suitably-revised notion of the overarching goal of the project of physics.

Over the last two decades, the quantum reconstruction program has generated intense interest in the quantum foundations community. Most of the attention thus far has focussed on the abstract quantum formalism (aQF), of which there are now several detailed reconstructions. The other parts of the formalism (QSA, SSC, QRs, QWEs)—see Fig. 2—have, in contrast, received relatively little attention. Meanwhile, the *interpretation* of quantum reconstructions is in the early stages.

But, before going further, we will first consider in more detail reconstruction as a methodology *per se*. In particular, we examine how the reconstructive methodology has been used to elucidate the theories of classical physics as a way of better understanding just why reconstruction of quantum theory is both a natural part of the life-cycle of the theory and is prerequisite for its proper interpretation. We shall then return to the reconstruction of quantum theory (Sec. V), and survey some of the recent interpretational insights that have been extracted from such reconstructions (Sec. VI).

IV. THE METHODOLOGY OF RECONSTRUCTION

A physical theory must somehow balance two very different demands. It must not only allow us to better *grasp* some aspect of the workings of the physical world, but it must yield a *workable* conceptual and mathematical tool that allows us to describe actual laboratory experiments and make precise predictions that conform to the brute facts of experience. In the development of a theory, if push comes to shove, the demand for workability tends to win out—rather as an individual tends to opt for safety over curiosity or self-actualization (as per Maslow’s hierarchy of needs), a theory’s survival in the scientific community normally depends far more on its capacity to successfully grapple with empirical regularities in the domain of interest and on the economy and usability of its mathematical formalism than on its intuitive graspability.

As a consequence, a freshly developed physical theory is inevitably a *compromise*. This tends to manifest in two main ways. First, there will be features of the theory’s mathematical formalism that (i) have been compelled to some degree by specific phenomena rather than being an expression of a more general physical idea or principle; (ii) are a compromise between conflicting physical ideas or desiderata; or (iii) are the product of somewhat non-physical rationales (such as ‘mathematical simplicity’). Second, the theory may refer to physical entities that are not directly observed (and are not directly observable according to the theory), but which form an essential part of the theory’s conceptual scaffolding.

Accordingly, once a new physical theory has been developed (the *developmental* phase) and its empirical power has been sufficiently demonstrated (the *proving* phase), there typically follows a *reflective phase* in which attempts are made to rectify these perceived weaknesses. In this reflective phase, *reconstruction* refers to the methodology whereby one elucidates the physical meaning and origin of physically obscure aspects of the mathematical formalism by deriving these from more fundamental or intuitively graspable physical ideas or assumptions. In the reflective phase, attempts are also made to eliminate the need for any unobservable entities by tracing the relevant parts of the formalism back to what can be directly observed, and such efforts are often part of—or directly or indirectly lead to—reconstructive

work.

A. Reconstruction in Classical Mechanics

Although Newtonian mechanics is generally regarded as a near-ideal theory (especially in contrast to quantum theory), its reconstructive phase witnessed a series of efforts to address its perceived formal and conceptual defects. For example, Newton’s framework brings up the following questions:

1. *Momentum as quantity of motion.* Why is the quantity of motion associated with a body $m\mathbf{v}$? In particular, why not a scalar quantity such as Descartes’ mv , or some more general vector $\mathbf{f}(m, \mathbf{v})$ or scalar quantity $g(m, \mathbf{v})$?
2. *Equation of motion.* Why is $\mathbf{F} = d\mathbf{p}/dt$ the equation of motion?
3. *Parallelogram of forces.* Why do the forces that act on a single body combine vectorially?

In each case, reconstructive work—some of which continues to this day—provides answers to these questions, and thereby grounds these crucial aspects of the Newtonian framework in more fundamental physical ideas and principles. The first and third of these are considered in some detail in Sec. IV A 1 and IV A 2 below. For the second, see (Darrigol 2019) and (Goyal 2020, §2.3).

The Newtonian framework also posits entities that are not directly observable, namely absolute space, absolute time, and the notion of force. Reconstructive work has not been particularly successful in obviating the need for these, but has spurred important advances. Mach argued that the notions of absolute rotation and absolute acceleration relative to an unobservable absolute space should be abandoned in favour of the idea of rotation and acceleration relative to the (observable) ‘fixed stars’. This notion—dubbed ‘Mach’s principle’ by Einstein—was a key input to Einstein’s general theory of relativity, but as that theory recovers Newtonian mechanics in a limiting case in an empty universe, it cannot be regarded as successfully implementing Mach’s principle. Poincaré noted the disconnect between the notion of absolute time and the conventionality inherent in any experimental determination of the simultaneity of distant events, which was likely a key input to Einstein’s special theory of relativity (Darrigol 2004, 2005). Finally, the notion of force was criticized by Hertz (Lützen 2005, Ch 4), who attempted to reconstruct mechanics without making use of this notion (Hertz 1899). However, his reconstruction is unsuccessful insofar as it introduces new and rather abstract unobservables (hidden cyclic systems) of its own (Lützen 2005, Ch. 6).

1. *Momentum as quantity of motion*

Descartes’ notion that the motion of a system conserves its total scalar ‘quantity of motion’ was a key input to the development of classical mechanics, in large part owing to its intuitive graspability¹⁰. Yet, in the subsequent

¹⁰ Descartes’ rationalization: ‘It is obvious that when God first created the world, He not only moved its parts in various ways, but also simultaneously caused some of the parts to push others and to transfer their motion to these others. So in now maintaining the world by the same

development of the laws of collinear collisions, this notion quickly came into conflict with two physical desiderata, namely that:

- (i) the total quantity of motion be conserved not only asymptotically (before and after the collision) but also at the moment of collision itself (i.e. *continuous* conservation, rather than merely asymptotic conservation); and
- (ii) the total quantity of motion be conserved during an inelastic collision.

As a result, Descartes' principle morphed¹¹ into a new conservation principle, the conservation of momentum, with momentum being a *vectorial* quantity of motion.

Thus, in retrospect, one can regard the conservation of momentum as a *compromise* between an intuitively attractive idea (Descartes' conservation principle) and other physically-motivated desiderata (applicability to inelastic collisions; continuous conservation). The principle of momentum conservation is mathematically expressible and workable, but is not faithful to Descartes' original idea—a universe can 'wind down' and yet conserve total momentum¹². In Newton's framework, it is rationalized via the idea that forces occur in opposed pairs ('action and reaction is equal and opposite'), an idea markedly different from Descartes'.

Descartes originally posited that a body's scalar quantity of motion had the mathematical form mv , which Newton (and others) vectorized to give momentum as $m\mathbf{v}$. The only motivation for these expressions appears to have been mathematical simplicity, leaving open the question of whether other mathematical expressions were physically viable.

Systematic derivations of the quantities of motion began to appear in the early twentieth century, apparently spurred by revisions to dynamics forced by special relativity, and have continued to appear ever since¹³. For example, in one recent derivation, the mathematical form of kinetic energy is derived from (i) a specific highly symmetric elastic collision, (ii) the requirement that some total (additive) scalar quantity of motion is asymptotically conserved; and (iii) Galileo's principle of relativity (Goyal 2020, §2). This derivation also generalizes to the special relativistic case, yielding the correct relativistic energy (Goyal 2020, §3). Moreover, an earlier argument by Schutz shows that asymptotic conservation of total kinetic energy and relativity together imply that total *momentum* is also asymptotically conserved (Schütz 1897).

Thus, reconstruction reveals the intimate relations between the scalar and vectorial conservation principles, and shows that the mathematical forms of the corresponding quantities of motion follow from the notion of conservation once supplemented by a fundamental *kinematical* symmetry principle (Galileo's principle of relativity, O4).

action and with the same laws with which He created it, He conserves motion; not always contained in the same parts of matter, but transferred from some parts to others depending on the ways in which they come in contact.' (Descartes 1982, II.42). Descartes' conservation principle is a symmetry principle, specifically an eliminative principle (Goyal 2020, §5.1.1)

¹¹ For more historical detail on this process of transformation, see (Goyal 2020, §5.3.1).

¹² Newton (amongst others) asserted that atoms were hard bodies that collide completely *inelastically* (Scott 1970, pp. 4–5). Hence the fundamental importance of formulating laws applicable to inelastic collisions.

¹³ For a sample of such derivations, see (Goyal 2020, §4).

2. Newton's parallelogram of forces

In his *Principia*, Newton obtained his so-called parallelogram law for combining forces that act on a singly body by considering the changes in the body's motion that results in the special case where the forces act impulsively at separate moments of time, and then combining the resultant changes. The argument thus depends upon Newton's second law that describes the effect that a force has upon the motion of a body, a law whose deeper origin is itself in question. And since the argument considers the case where the forces are unbalanced (and so cause a change in motion), it does not cover the case where the forces are in static equilibrium.

Following Newton's formulation, numerous attempts were made to place the parallelogram of forces on a sounder conceptual footing (Lange 2011). One of the most incisive arguments is due to d'Alembert (and subsequently refined by many others), which shows that the parallelogram law can be derived largely from elementary symmetry assumptions¹⁴:

1. the resultant of two parallel forces has magnitude equal to the sum of the magnitudes of these forces, and points in the same direction.
2. the resultant of a number of forces is commutative and associative.
3. the resultant of two forces is rotationally covariant.
4. the resultant of two equal forces varies continuously with the angle between these forces.

The first assumption is a particular case: if two forces are parallel, their magnitudes simply add. As it happens, one can simplify this further: additivity of magnitudes follows from the requirement that parallel-force composition is associative and continuous at a point¹⁵. The second and third assumptions are *symmetry* requirements, the second appealing to compositional symmetries, the third implicitly appealing to the isotropy of space (C_3). Finally, the fourth assumption appeals to the requirement of continuity—the resultant of two forces should change *gradually* as the forces gradually change.

As this argument does not depend upon any relation between force and motion, it applies equally in the static and dynamic cases, and is independent of any law connecting force and motion. Moreover, it shows that the parallelogram law follows very generally from compositional and spatial symmetries, together with the primitive idea that the magnitude of a force can be quantified by a real number. As one can see, very little of the intuitive notion of force is left: the assumptions essentially only require that force is something that has a direction and real-valued magnitude and can be composed with other forces in a manner that satisfies basic compositional symmetries.

¹⁴ For mathematical details, see, for instance, (Aczél and Dhombres 1989, Ch. 1).

¹⁵ See, for instance, (Aczél 1966, §6.2).

B. Reconstruction in Electromagnetism

Reconstructive work in classical mechanics has generally served to more solidly ground aspects of the Newtonian mathematical framework in more elementary physical principles, but has not fundamentally challenged the classical conceptual framework in which it is embedded. However, in electromagnetism, the reconstructive method historically led to a profound transformation in the interpretation of the theory—and of the nature of space and time—which was wholly unexpected.

1. Faraday's field conception

Beginning in the early nineteenth century, theories of electromagnetic phenomena developed along two parallel streams. The major stream sought to embed electric and magnetic phenomena within the existing Newtonian framework by ascribing a new time-independent property—charge—to each particle, and by formulating new force laws (Coulomb, Ampère, Biot-Savart, Weber, etc.), patterned after Newton's law of gravitation, that govern the interaction between static and moving charges and between current elements. The minor stream, initiated and sustained by Faraday, sought to understand electromagnetic phenomena through a novel conception of space-filling electric and magnetic *fields* produced by—and influencing—charges and currents.

A rather crucial feature of both streams is that their key features were strongly shaped by the peculiarities of the phenomena of interest, rather than by some general *a priori* principles as was the case with classical mechanics. Thus, in Faraday's conception, the need for *two* distinct fields (electric, magnetic) as opposed to just one; the fact that these fields interact in the specific way that they do with charges and each other; and the fact that charges (electric monopoles) exist but magnetic monopoles do not—all were ultimately justified by the need to bring order to the phenomena of interest. In contrast, as noted previously, classical mechanics was largely based on general principles—the constant velocity of an isolated particle; the physical equivalence of all inertial frames; conservation principles—which seemed to spring more from the instinct to bring an ideal mathematical order to the physical realm rather than from an attempt to make sense of any specific regularities in the phenomena of interest.

2. Maxwell's equations and their interpretation

Maxwell's equations (1865), the heart of electrodynamics, are in essence a mathematical clothing of Faraday's field conception (although Maxwellian electrodynamics incorporates important elements from the force-law stream, for example in the form of Lorentz's force law). Due to the predicted existence of electromagnetic waves, and the belief that such waves—like sound waves—require a *medium* for their propagation, it was almost universally assumed until at least the end of the century that Maxwell's equations apply to a *particular* inertial frame, namely the frame that carries the requisite medium ('aether'). This naturally accounted for the fact that, unlike Newton's equations of motion, Maxwell's equations are not invariant under Galilean transformations. However, this 'natural' interpretation generated a clear tension with Galileo's principle of relativity (O₄) and with the Newtonian mechanics that

incorporated Galileo's relativity principle.

Meanwhile, Maxwell's equations possess a rich mathematical structure, which was gradually brought to light. In particular, as discovered by Lorentz in 1892 (although to some degree anticipated by Voigt in 1887), these equations are invariant under the so-called Lorentz transformations, which suggests that these equations *are* valid in non-aether inertial frames provided that one introduce new abstract frame-dependent space and time coordinates. However, the physical meaning of these transformations and these abstract coordinates was unclear. Moreover, the aether hypothesis also faced a mounting challenge due to (i) the difficulty of coming up with a single aether model capable of supporting the full range of physical phenomena (such as the apparently frictionless passage of planetary bodies through the aether, as well as the propagation of light at such high speeds); and (ii) the failure of experimental attempts to detect motion relative to the aether (in particular the 1887 Michelson–Morley null result).

3. *Einstein's reconstruction and its interpretational implications*

The meaning of the Lorentz transformations, and of the failure to detect motion relative to the aether, was dramatically elucidated in 1905 by Einstein, who *reconstructed* the Lorentz transformations in an operational framework. Astonishingly, he derived the transformations without any direct reference to Maxwellian electrodynamics, the field concept, or the aether hypothesis. Instead, the Lorentz transformations were traced back to elementary spatial and temporal measurements carried out using rods and clocks. The derivation was based on two key assumptions: (i) the one-way speed of light is independent of the speed of the source (a reasonable extrapolation of experimental facts to date); and (ii) Galileo's principle of relativity. Einstein reconciled these two seemingly contradictory ideas by positing a reasonable definition of light-based synchronization of distant clocks, and then showed how they led to the Lorentz transformations. He went on to build a new mechanics compatible with the Lorentz transformations, which led to specific testable predictions.

Einstein's reconstruction thus showed that the Lorentz transformations *were* compatible with Galileo's principle of relativity, but that there was a cost: observers in different inertial frames would, in general, disagree as to the spatial distance and temporal duration between two events. Thus, even if there were an absolute distance and duration between the events (as per the Newtonian classical conceptual framework—see Sec. II A 2), it would be experimentally inaccessible.

Einstein's reconstruction was so compelling, and the various solutions hitherto proposed to address the challenges faced by the aether hypothesis sufficiently unattractive, that the aether hypothesis was effectively abandoned within a few years, and a new interpretation of Maxwell's electrodynamics established. Moreover, the very notion of absolute space and time was brought into question by the reconstruction's implication that absolute distances and durations (if they exist) are experimentally inaccessible, and it became the norm to require that a theory or model be Lorentz invariant.

V. RECONSTRUCTION OF QUANTUM THEORY

The development of quantum theory bears many important similarities to that of electrodynamics. In the absence of general physical principles of sufficient power, both theories were strongly shaped by the peculiarities of—and observed regularities in—the specific phenomena of interest. And the mature formalisms of both theories possess mathematical features and mathematical structure whose deeper physical meaning was initially obscure. As we have seen above, the physical meaning of the principal symmetry of Maxwell’s equations, namely the Lorentz transformations, was elucidated by Einstein’s reconstruction.

Can the reconstructive method similarly elucidate the various striking mathematical features of the quantum formalism? For example, why does the formalism employ complex numbers and why are complex numbers so well suited to the expression of the theory? Why are dynamics represented by unitary transformations, rather than a broader class of transformations? Why does the Born rule (which connects complex-valued states and outcome probabilities) take the specific mathematical form that it does rather than a more general form? Why is the tensor product operation appropriate to construct the states of a composite system? Why are the states of identical particles subject to the symmetrization postulate? A reconstruction of the quantum formalism which traces these abstract mathematical features to clearly-stated physical principles in an operational framework—analogue to Einstein’s reconstruction of the Lorentz transformations—could yield deep, unexpected insights into the reality that is so astonishing well described by quantum theory.

Recognition of the importance of reconstruction for elucidating the quantum formalism was not lost on the founders. For example, Heisenberg recognized that it would be highly desirable if the quantum formalism could somehow be derived using his uncertainty principle as a key axiom. And, in his 1946 Nobel lecture (Pauli 1998), Pauli expressed the view that his exclusion principle (even after incorporation into the symmetrization postulate) called for a deeper explanation in the form of a ‘rigorous derivation’ from more general assumptions, but that no such explanation had hitherto been forthcoming¹⁶.

Broadly speaking, reconstructive attempts prior to the 1980s tended towards abstract, intricate systems of axioms, which made little impact (Grinbaum 2007a,b). Since the 1980s, the program of reconstruction has gained fresh impetus and inspiration from several directions, and has gradually gained traction in the foundations of quantum physics community¹⁷. There now exist a number of fairly rigorous reconstructions of the abstract quantum formalism

¹⁶ ‘Already in my original paper I stressed the circumstance that I was unable to give a logical reason for the exclusion principle or to deduce it from more general assumptions. I had always the feeling and I still have it today, that this is a deficiency. Of course in the beginning I hoped that the new quantum mechanics, with the help of which it was possible to deduce so many half-empirical formal rules in use at that time, will also rigorously deduce the exclusion principle. Instead of it there was for electrons still an exclusion: not of particular states any longer, but of whole classes of states, namely the exclusion of all classes different from the antisymmetrical one.’ (Pauli 1998, p. 32).

¹⁷ A comprehensive overview of the reconstruction program, which details the full range of approaches and their inspirations, assumptions, and techniques, has yet to be written. This is not surprising: the diversity in the approaches’ conceptual starting points and mathematical techniques is immense, and the program is still in flow. Nevertheless, various perspectives—some written by researchers in the field and some by philosophers of physics—do exist. For example, in (Hardy 2013, §2), Hardy—one of the first to present a compelling reconstruction of the aQF—offers brief but insightful remarks on the history of reconstruction. Grinbaum (Grinbaum 2007a,b), a philosopher of physics, provides another interesting perspective.

in finite dimensions, and a fewer number of reconstructions of many of the other parts of the formalism¹⁸.

A. Informational perspective

One of the major forces behind the renewed interest in reconstruction has been the *informational* perspective on physical theory¹⁹, which is aptly summarized in Wheeler’s slogan ‘*It from Bit*’ (Wheeler 1989, 1990):

“ ‘*It from bit*’ symbolizes the idea that every item of the physical world has at bottom—at a very deep bottom, in most instances—an immaterial source and explanation; that which we call reality arises in the last analysis from the posing of yes-no questions and the registering of equipment-evoked responses; in short, that all things physical are information-theoretic in origin, and this in a participatory universe.”

and

“What we call reality consists of a few iron posts of observation between which we fill an elaborate papier-mâché of imagination and theory.”

These views colourfully echo Mach’s view of physical theory as first and foremost an economical representation of physical observations. But they go significantly further. First, the assertion that the ‘physical’ world has an ‘immaterial source’ places *observation* at the centre and decisively demotes the notion of matter to part of the elaborate conceptual papier-mâché that we build in order to make sense of these observations²⁰. Second, Wheeler speaks of measurement not as a passive observation of that which exists, but rather as an *active* process whereby we *pose questions* to nature and *evoke a response*. This view of measurement arises from a particular reading of the quantum formalism, namely one that takes measurement as a primitive, and moreover one that posits that measurement is not a passive registration of pre-existing properties but rather that measurement outcomes are in some sense *co-created* through the act of measurement.

The fertility of the informational viewpoint has been strongly supported by the emergence of the fields of quantum information and quantum computation. These fields’ many successes and striking discoveries (such as the possibility of secure information transfer and better-than-classical computation) have deepened the conviction that quantum theory rewards an informational perspective, in particular that it provides a new lens through which to look at the quantum realm which may well allow us to make new progress in understanding the nature of quantum reality.

¹⁸ A range of talks from leading researchers in the field presented at the ‘Reconstructing Quantum Theory’ workshop held at Perimeter Institute in 2009 are available online at <https://pirsa.org/C09016>.

¹⁹ The informational perspective in physics precedes quantum theory, and can be traced at least as far back as the development of statistical mechanics. For a more detailed treatment of its development, see (Goyal 2012).

²⁰ In contrast, in his *Science of Mechanics*, Mach recognises that observations are central and casts doubt on the Newtonian notion of absolute space due to its lack of direct observability, but does not go so far as to question the notion of matter as primitive (Mach 1919).

B. Operational Framework

Most reconstructions take place in an *operational framework*. In essence, this is an idealized, abstract representation of the laboratory workbench, the place where theory comes into contact with physical reality.

I. Experimental set-up.

The operational framework rests on the idea that a *physical system* is subject to an experiment, namely a sequence of *measurements* and *interactions*. The measurements and interactions are presumed to be implemented by objects that are, at least to some extent, classically-describable (*i.e.* ‘macroscopic objects’), and whose settings are under an experimenter’s control. Measurements differ from interactions in that measurements yield macroscopically-observable events (such as flashes on a scintillation screen), and it is through these events that the physical system comes to be indirectly known. Idealized measurements are typically assumed to be *repeatable* in the sense that immediate repetition of a measurement on a physical system yields the same outcome with certainty.

i. Notion of physical system. The idea that an experiment is performed upon a ‘physical system’ consists of two distinct notions. First, that there exists some entity that *persists* for the duration of the experiment. That is, in spite of interactions and measurements performed upon it, there is some meaningful sense in which there is something which retains its identity over the course of the experiment, so that one can say that a measurement at t_2 is performed on ‘the same’ system as the measurement at t_1 . The notion of persistence thus provides a minimal yet crucial means to link outcomes obtained at different times.

Second, that there is some definite sense in which all of these interactions and measurements are probing the same *aspect* of the physical object. We accordingly say that the experiment is performed upon a *physical system* or abstract physical object. In experiments on physical objects deemed to be classically-describable, this notion is usually implicit—we understand that experiments on real objects (such as billiard balls) must be carefully circumscribed in order to extract meaningful information. For example, consider an experiment which is designed to probe the spatial behaviour of a classically-describable billiard ball. A preparation would consist in fixing the initial position and velocity of the ball. Subsequent measurements would need to be restricted to those which only probe these degrees of freedom (position, velocity)—a measurement whose outcome depends upon the ball’s rotational motion or, say, its material composition or colour, would need to be excluded. Similarly, interactions with the ball would need to be restricted to those that do not couple the spatial (position, velocity) degrees of freedom and its non-spatial degrees of freedom. For instance, interactions which change the ball’s velocity in a way that depends upon its colour would need to be excluded. With these restrictions in place, the outcomes of measurements on the ball would be independent of pre-preparation interactions with it (such as re-painting it a different colour), and the experiment would be probing only the spatial sub-component of the actual physical object. One could then abstractly describe this sub-component as ‘a particle’, namely, an abstract object whose only time-dependent properties are position and velocity.

As described in (Goyal 2008, § II A), these ideas can be generalized in a way that can be applied to non-classical

systems—specifically systems to which *a priori* we (i) cannot attribute properties and (ii) cannot assume these properties can be passively observed—via the notion of *closure*. The closure condition can be used to operationally establish that two measurements are probing the same aspect of an object. One can then operationally define the *set* of all measurements that probe the same aspect of the object, and define a set of interactions that can be thought to act wholly on the abstract object thereby defined. In short, the ‘physical system’ refers to some *operationally-defined aspect* of a physical object.

ii. Nature of the agent. In the operational framework, it is usually left implicit that there exists an entity—an *agent*—that is capable of *observing* or *registering* measurement outcomes, and that the entity possesses the capacity to make changes in device settings without being influenced by the system under study or by past or future measurement outcomes. These notions are usually implicit in experimental science, but the precision required in reconstructive work is such that they sometimes need to be made explicit.

2. *Abstraction of key notions.*

The operational framework is an *abstraction* of the laboratory workbench, not a *representation* of it. Accordingly, a great deal that one tends to unconsciously associate with an ‘experiment’ is not an intrinsic part of the operational framework, and indeed turns out not to be needed in many reconstructions of the abstract quantum formalism.

For example, we tend to think of an agent as an embodied being localized in space; a physical system as an object that is spatially localized in our laboratory at all times; or a measurement as carried out by a chunk of equipment in one corner of a laboratory. But the operational framework abstracts away all of these *spatial* notions. So, a *physical system* is simply an entity that *persists*—it does not necessary exist *anywhere* in particular at a given moment in time. A *measurement* is an abstract parameterized process that acts on a physical system to generate an *outcome* and to output the same physical system—it is not a spatially localized piece of equipment. The agent is simply *an entity* that exists and persists over time, and is capable of *observing outcomes* and of *freely acting* to change *settings* associated with measurement and interaction devices—it is not a spatially localized human being.

On the other hand, the notion of *temporal order* is essential to the operational framework. It is assumed that the measurements and interactions occur in a well-defined temporal sequence. In particular, temporal order is essential to the notion of closure (which is essential to the definition of a physical system, and the set of measurements and interactions), and the notion of ‘immediately afterwards’ is required to ground the above-mentioned notion of repeatability.

3. *Classical component of reality as a portal.*

The operational framework makes essential reference to physical objects that an experimenter can observe and manipulate. In so doing, the framework presumes that there is a component of physical reality that is well described by our everyday object model of sensations and/or by classical physics. For example, it is assumed that the apparatus

and the agent both persist over time; that the agent can passively observe measurement outcomes; that the apparatus has well-defined properties (in practice, such things as dial-positions and knob settings) that can be unproblematically adjusted by the agent.

This *classical component* of physical reality then effectively acts as a *portal* through which we probe some other part of physical reality, be it classically-describable or not. For example, in an electron diffraction experiment, the experimenter has access to a device that generates the accelerating voltage and to devices that generate deflecting electric and magnetic fields; and can observe flashes on a scintillation screen which he interprets as evidence of electronic collisions. But the behaviour of the ‘electron’—the microphysical object which is commonly supposed to underpin the scintillations—is not classically-describable. But it is only *through* these classically-describable bits of equipment that we learn about its behaviour.

To be clear: we cannot *directly* access the non-classical because we do not know if our classical notion of physical object (namely, an persistent entity that possesses properties) is valid, and we do not know if we can passively observe the properties of these objects. But these notions are the basis for defining ‘an experiment’, and underpin the possibility of carrying out many trials of ‘the same’ experiment and accumulating statistically-analysable data.

For instance, when we speak of a measurement device having a *setting*, we are evoking an object–property model, in particular attributing the setting to the object itself. And we also presume that an agent can passively observe that setting. These are all ‘common sense’ notions drawn from everyday experience with physical objects, assumptions that are enshrined in the classical conceptual framework. It is far from clear what is left of the notion of ‘experiment’ if such assumptions cannot be made.

In short, it would seem that access—at least *scientific* access—to the non-classical component of reality requires a classical portal. This is not, of course, to say that one cannot access the non-classical component of reality without, for example, the possibility of changing ‘settings’ in a reliable way, or without the possibility of carrying out multiple trials of an ‘experiment’. Indeed, such *non-scientific* access to non-classical reality might well correspond to much of so-called subjective experience.

C. Physical principles

Once the operational framework is in place, one must posit *physical principles* which precisely articulate guesses or hunches about the nonclassical physical reality which is manifested in experiments. It is characteristic of most recent operational reconstructions that these physical principles refer primarily to the *data* gathered in experiments, rather than attempting to *directly* posit features of the physical system under scrutiny. For that reason, such reconstructions are often referred to as *informational* or *information-theoretic* reconstructions, with the latter particularly common if the machinery of Shannon’s information theory or Bayesian inference is employed.

I. Wootters' derivation of Malus' law

An excellent illustration of the kinds of physical principle that are employed in recent reconstructive work is afforded by one of the earliest informational reconstructive results due to Wootters (Wootters 1980). Consider an experiment on a physical system subject to measurements that yield one of *two* possible outcomes (which we label 1 and 2), where each measurement is parameterized by a single angle²¹. This is an abstraction of Stern-Gerlach measurements performed on spin-1/2 systems.

Consider a game in which Alice *prepares* one such two-outcome system by performing a measurement with θ , and then transmitting that system to Bob if the outcome happens to be 1 (she discards those systems that yield outcome 2). On receipt, Bob performs a measurement with setting $\theta' = 0$, and records its outcome (either 1 or 2). In total, Alice sends n identically-prepared systems to Bob.

Now, let us suppose that reality is lawful in the minimal sense that the *probability* of Bob's outcome is determined. That is, the probability, p , of his obtaining outcome 1 is a function of θ . It follows that, if Bob were to know the function $p(\theta)$, then he could make a reasonable guess about the setting angle, θ , given the frequency with which he obtains outcome 1. That is, the string of outcomes that he obtains in the n measurements provides him with some information about θ . Accordingly, one can view this as a game in which Alice imperfectly transmits the angle θ to Bob by 'encoding' the angle in n two-outcome physical systems.

Wootters now posits that the laws of nature are such that the amount of information that is transmitted—quantified using the relative entropy—is *maximized* in the limit as $n \rightarrow \infty$. Wootters then effectively shows²² that, if one assumes that the prior probability over θ , $\Pr(\theta | I)$, is uniform, then $p(\theta) = \cos^2(m(\theta - \theta_0)/2)$, where $m \in \mathbb{Z}$. This *agrees* with the predictions of quantum theory for a two-outcome system, and is known as Malus' law.

D. Informational reconstructions

Wootters' derivation of Malus' law is minimalistic in its assumptions, and is a powerful illustration of the potential of informational approaches to quantum theory. However, this minimalism does not survive the passage to reconstructions of the quantum formalism proper.

In particular, most reconstructions introduce the notion of *state* in the form of a mathematical object that is associated with the physical system at each moment in time, and whose role is defined operationally as that which allows the prediction of the outcome probabilities of any²³ measurement that could be performed on the system. In contrast, in the derivation of Feynman's formulation of quantum theory (Goyal et al. 2010), one considers a *transition* between given outcomes of successive measurements and associates a mathematical object—a pair of real numbers,

²¹ It is understood that these measurements satisfy the conditions of the operational framework. In particular, these measurements are repeatable, and satisfy the closure condition.

²² For more mathematical details, see (Goyal 2012, §4.1).

²³ As per the operational framework (Sec. V B i i), the set of possible measurements needs to be circumscribed by a specific procedure.

which eventually becomes a complex-valued *amplitude*—to that transition, one role of which is to determine the transition probability.

The introduction of the notion of state or transition amplitude adds a layer of abstraction to the reconstruction, which is then reflected in the postulates to some degree. For example, Hardy (Hardy 2001a,b) formulates postulates that refer to the number of degrees of freedom associated with a state. Similarly, the derivation of Feynman’s rules emerge as a pair-valued quantification of an experimental logic. The interpretation of this abstraction layer poses a considerable challenge.

Nonetheless, certain specific assumptions which are then made in order to give shape to the resulting mathematical structure can be more readily understood. One class of postulates concerns bipartite systems. For example, one operationally-expressible postulate known as *tomographic locality* posits that the state of a bipartite system can always be determined from the statistics of a sufficient number of different joint measurements performed separately on the two sub-systems. This postulate is employed to good effect by Barrett (Barrett 2006) to account for at least some of the structure of the quantum formalism, such as the tensor product rule for determining the state of a composite system when its subsystems are in known pure states. Hardy interprets one of his postulates as an expression of tomographic locality (Hardy 2013).

Another class of postulates concern the behavior of individual systems. For example, I have shown that it is possible to reconstruct the aQF via the Feynman rules of quantum theory by suitably formalizing the notion of complementarity and by introducing a *no-disturbance* postulate which posits that certain measurements which yield no useful information about a physical system also do not disturb its state in any detectable way (Goyal 2014, Goyal and Knuth 2011, Goyal et al. 2010). Remarkably, the complex structure of the quantum formalism, together with the rule that determine the outcome probabilities, are completely determined.

Here, the no-disturbance postulate is an expression of the idea that measurement is an *active* process. In particular, that the acquisition of information about a system in general forces a change in the state of the system, and that, in the limiting case that the measurement yields *no* information about the system (it simply tells us that the system exists), no such change occurs.

VI. INTERPRETATION OF QUANTUM RECONSTRUCTIONS

As described in Sec. III D 2, an ideal interpretation of quantum reconstructions would yield (i) a quantum conceptual framework (precisely articulated in a form analogous to the classical conceptual framework as in Fig. 1); and (ii) a revised notion of the project of physics that ties together the framework’s assumptions.

The interpretation of quantum reconstructions is presently in its infancy. Be that as it may, some general interpretative implications of reconstructive work are already clearly visible, which I shall summarize in Sec. VI A below. More specific implications are also starting to come into view. In particular, over the last few years, I have developed an interpretation of the quantum symmetrization algorithm based on a reconstruction of the same (Goyal 2015), which gives rise to a new understanding of the nature of identical quantum particles (Goyal 2019, 2022). The

key ideas are summarized in Sec. VI B. An interpretation of the reconstruction of Feynman’s rules is also underway, which I believe will yield a metaphysically sharp formulation of Bohr’s principle of complementarity.

A. Interpretative implications from reconstruction of the abstract quantum formalism

The abstract quantum formalism (aQF) is the core of quantum theory. Although an abstract shell rather than an explicit model of a particular physical system, it is highly contentful. For example, the striking protocols in quantum information (such as quantum teleportation or quantum cryptography) or schemes for quantum computation only make use of the aQF. Key no-go theorems, such as Bell’s theorem and the Kochen-Specker theorem, can likewise be formulated wholly within the aQF. And more broadly, most of the familiar non-classical features of quantum theory, such as the existence of entangled states, the change of quantum states due to measurement, and the notion of complementarity (in a variety of forms), can be articulated entirely within the aQF.

For these reasons, the aQF is the prime reconstructive target, and its interpretation is of the greatest importance. Below we make some preliminary remarks on some of the general interpretative insights that can be drawn from many of the operational reconstructions of the aQF (such as Chiribella et al. (2011), Goyal (2014), Goyal et al. (2010), Hardy (2001a)).

1. Notion of space.

The operational reconstructions of the abstract quantum formalism make no explicit reference to the notion of space—to the idea that objects are localized in space, or that space has a certain dimension, topology, or metric. In its stead, the operational framework presumes that an observer has some ground for saying that *this* measurement is performed on *that* system, and *this* outcome (rather than some other) is obtained.

In practice, an experimenter grounds such assertions on spatial perception, and relies upon the assumed persistence and continuous locations of macrophysical objects. But reconstruction of the aQF makes clear that we do not *need* the notion of space (with all of the rich set of ideas that accompany it) *per se*—the observer only needs *some* means to ground such assertions.

2. Notion of measurement.

As indicated in Sec. V B I, whereas one ordinarily thinks of a measurement as implemented by a physical device localized in space, the operational framework abstracts from the laboratory workbench to the extent that all that is left of the notion of *measurement* is that it is an abstract parameterized physical process that acts on the physical system and yields an *outcome*. This more abstract notion of measurement is indeed in keeping with the above remarks on space.

This more abstract notion of measurement is, in fact, indispensable in the usual applications of the quantum formalism. For example, one commonly regards the outcomes of several localized measurements as a single outcome

of a *joint* measurement, and that joint measurement is not located anywhere in particular. At a more fundamental level, a more abstract notion of measurement frees one's imagination to consider processes of actualization—such as Penrose's gravitationally-induced objective reduction (OR) or GRW's collapse model—which are not tied to macroscopic measurement devices.

3. *Notion of time.*

In contrast to space, the notion of *temporal order* is central to the operational framework. In particular, the observer *must* be able to say '*this* measurement was performed, yielding such-and-such outcome; *then* that measurement was performed...', and so on. Thus, the abstract quantum formalism depends on the notion of time, at least in the very specific sense that it is assumed that an observer has the means to temporally order their experiences. In addition, the notion of 'immediately afterwards' is required to ground the notion that a measurement is repeatable.

4. *Primacy of time over space.*

Based on the above remarks about space and time in the operational framework, the abstract quantum formalism—the core of quantum theory—does not 'know' about space in the Newtonian sense. It does, however, 'know' about time in the limited sense of temporal order and immediate succession. This suggests an obvious interpretation, namely that the aQF describes a reality that *precedes* space²⁴. This, coupled with the fact that the aQF captures so much of what distinguishes quantum physics from classical physics, is intriguing on a number of levels.

First, the aQF was historically obtained via a process of abstraction from concrete quantum models designed to account for such phenomena as the spectral lines generated by excited atoms. Moreover, these models were—in the work of, say, Bohr, Heisenberg and Schoedinger—typically arrived at by subverting classical models of those same systems by introducing new ideas (such as quantization of angular momentum or wave-particle duality). But the aQF makes no reference to a gamut of fundamental classical notions such things as space, energy, and momentum. So, the fact that the aQF can be reconstructed without explicit reference to these classical notions is remarkable, and shows that the aQF is a free-standing structure that does not need to lean on the classical conceptual framework or on any particular classical physical theory (such as classical mechanics or electromagnetism).

Second, the idea of a richly-textured reality that precedes space may be scientifically useful. The idea is, in fact, quite prevalent in certain approaches to quantum gravity in the form of the posit that there exists some structure prior to space ('pre-space'), of which space is only an effective, approximate description. At a more metaphysical

²⁴ In this connection, it is useful to compare the aQF with its analogue in classical physics, to which one could refer as the abstract classical formalism (aCF). In the aCF, one speaks of an abstract system in state that evolves continuously, deterministically, and reversibly; of the existence of an ideal measurement whose outcome determines the state without disturbing the state; and of a system's subsystems always being in well-defined states which determine the state of the system. In the aQF, dynamics is—as in the aCF—reversible and deterministic, but the ideal measurement and compositional axioms differ.

level, as I describe in (Goyal 2019, 2022), the flexibility that this notion provides in conceiving of physical reality opens up new possibilities for conceptualizing identical particles (see Sec. VI B for a brief summary).

5. *State concept.*

In reconstructions of the aQF, the notion of state is introduced as a means of connecting together measurement outcome data. This is particularly clear in the Feynman pathway (Goyal et al. 2010), where amplitudes (initially *pairs* of real numbers) are introduced to connect together pairs of measurement outcomes obtained at different times, and where states are then built up as collections of these amplitudes (Goyal 2014). Thus, the mathematical state object reflects the chosen measurements as well as the physical system in question.

This contrasts sharply with the state concept in the classical framework, which is thought to describe the actual physical state of a system (such as the position and velocity of a particle) quite independently of whether or not it is observed, or indeed independently of whether there are any observers at all.

6. *Nature of the agent.*

As mentioned in Sec. V B 1, it is usually implicit in the operational framework that there exists an agent capable of (i) passively observing measurement outcomes, and (ii) changing measurement and interaction device settings ‘at will’. As described in Sec. II A 2, the agent also appears in this dual role in classical physics. However, the classical framework posits that, absent agential action, physical reality can be described and modelled without including the agent. The key interpretational question is then whether that remains the case in quantum theory, or whether there is a compelling reason to include the agent explicitly as part of the description of physical reality.

As described above, the notion of a quantum state is introduced as a means of connecting measurement outcomes in an operational context. That context is shaped by the actor-as-agent—the agent chooses the measurements to which to subject a physical system, and the set of quantum states that can be ascribed to the system is a function of that choice. As such, the quantum description of physical reality is—from the reconstructive point of view—inextricably context-dependent, and that context is shaped by agential action. This contrasts sharply with the situation in classical physics, wherein one is permitted to speak of the physical state of a system—or the universe as a whole—without regard to agential action.

We could then choose to say that quantum theory differs from classical theories in that it is a *context-dependent* formalism, namely that its formalism presupposes a classically-describable experimental context. In such a case, the agent (understood in the dual sense as above) does not play a role that is fundamentally different than in classical physics. Alternatively, one could assimilate the context to the agent (as appears to be favoured in QBism—see (Fuchs 2017), especially Fig. 1), so that the formalism is inextricably bound up with agential experiences and choices. My own view is that such an assimilation is unnecessary and is liable to lead to intractable difficulties, and that the first view—bearing in mind that the experimental context is agentially shaped—is the more fruitful.

Two final remarks. First, it is important to note that, insofar as the reconstruction of the aQF is concerned, the nature of the agent beyond that which is capable of serving a passive observer or an ‘at will’ actor are not specified. In particular, the notion of a bodily-extended, spatially-bounded agent is unnecessary—one could conceive of the agent as an incorporeal observer and actor. One need not even suppose that there are many distinct actors—one could conceive of a single (universal) actor.

Second, the agent-as-actor—namely the entity capable of actively changing experimental settings ‘at will’ (i.e. uninfluenced by the physical system or by future measurement outcomes)—is essential to the *carrying out* of actual experiments and its capabilities are implicit in the interpretation of experimental results. Although this is also true in classical physics, it is intriguing that certain key results in quantum foundations require that this be made explicit. For example, in the statement of Bell’s theorem, it must be made explicit that experimenters in each wing are ‘free’ in their choice of measurement settings, in particular that they are not in any way influenced by the present or previous state of the system upon which the measurements are being performed. It also must be made explicit that the state of the system at an earlier time is not influenced by what measurement settings will be chosen at a later time. This is sometimes referred to as the ‘no-conspiracy’ assumption, which is in fact implicit in virtually all experimental design.

B. The nature of identical quantum particles

Atomism seeks to account for our everyday experience of persistent objects by positing the existence of elementary entities that have continuous transtemporal existence. As described in Sec. II A 2, this metaphysical conception is incorporated into the conceptual framework of classical particle mechanics: macroscopic bodies are assumed to be composed of eternal (indefinitely persistent) point particles, each of which possesses time-independent properties (such as mass and charge) and moves continuously through space. Each such particle can in principle be reidentified, either by measurement of its distinct time-independent properties (if it is unique), or by tracking it precisely over time.

According to present day particle physics, ordinary matter is composed of elementary entities such as electrons and quarks (which, in turn, compose protons and neutrons). But, quantum theory raises the question as to whether—and to what extent—these ‘particles’ possess the various characteristics ascribed to them by classical physics.

In particular, the quantum formalism contains a specific algorithm—the quantum symmetrization algorithm (QSA)—which must be employed to model systems composed of *identical* particles, such as the two electrons in a helium atom. According to this algorithm, the set of allowable states of such a system are restricted to those that are either symmetric or antisymmetric, depending upon the type of particle in question. According to the standard interpretation of this restriction (known as the symmetrization postulate), there exists no measurement that can ‘address’ a particular particle in a system of identical particles. Consequently, particle reidentification is impossible. Such an implication, however, is directly at odds with the basic experimental data of particle physics: a bubble chamber images is said to show *particle tracks*, and the geometry of these tracks is used to determine these particles’ intrinsic properties. The question is, then, how to reconcile these two (theoretical and experimental) perspectives.

Over the last few years, I have developed an interpretation of the quantum symmetrization algorithm based on a reconstruction of the same (Goyal 2015), which reconciles these two perspectives (Goyal 2019, 2022). The essential idea that emerges is that, *contra* the atomistic conception, particles cannot be said to simply *exist*. Rather, the particle notion is simply an element of a model of detection events, and that model is only strictly valid in limiting cases. In general—as, for example, in a helium atom—a duality of object-models is needed to conceptualize the underlying microphysical reality. In one of these models (the so-called persistence model), there exist two distinct persistent entities. However, in the other model (the nonpersistence model), there exists but a single persistent entity—a holistic object, if you will—which manifests as multiple point-like events at a single moment in time. In general, *both* models must be mathematically synthesized in order to generate empirically accurate models.

This radical conclusion can be further illuminated through connection to metaphysical debates on such issues as event ontology, strong emergence, and monism *vs.* pluralism. For example, I argue that microphysical events have primacy over microphysical objects (a restricted form of event ontology); and that one can regard the formation of microphysical composites of identical particles (such as in the formation of a helium atom) as the emergence of a holistic object, and, conversely, the decomposition of such composites as the emergence of particles.

VII. CONCLUDING REMARKS

Quantum theory is our most successful physical theory, both in the sheer *range* of physical phenomena it is capable of describing and the *precision* with which it does so; and in its capacity to repeatedly make *unexpected novel predictions* that are subsequently borne out²⁵.

Quantum theory is also by far our most enigmatic physical theory. Despite a century of efforts to peer through the quantum veil, we are still left in the strange predicament of being in possession of an astonishingly powerful theoretical tool that we scarcely comprehend, *viz.* a tool that we do not know how to speak about in natural language in a coherent, comprehensive, and precise manner—let alone with any logical-philosophical precision and systemacity—without resorting to theoretical jargon or leaning on mathematical crutches.

The elucidation of quantum theory is of vital importance for the development of physics, in particular for development of theories of quantum gravity, and—as the development of the fields of quantum information and computation attests—for the full-bodied exploration and technological harnessing of the quantum realm. But, the importance of such elucidation goes far deeper: quantum theory forces reconsideration of a world-view that has influenced or guided the development of Western society for three centuries. Hence, a vivid understanding of quantum theory, one precise and comprehensive enough to be put alongside the mechanical conception of physical reality, would likely have widespread consequences for many compartments of human culture and human knowledge.

²⁵ These include the prediction of Bell-violating correlations, which are generally interpreted as a manifestation of nonlocality. According to Lakatos, novel predictions that are unexpected relative to previously-existing theories, and which are subsequently experimentally verified, are a hallmark of progressive research programs (Lakatos 1978). In that respect, it is astonishing that quantum theory has continued to form the heart of a progressive research program for a century, and shows no signs of flailing. In comparison, so-called regressive research programs typically play catch-up, merely accommodating new findings (whether they be experimentally discovered or anticipated by rival research programs).

I have argued that it is entirely reasonable to expect to be able to understand quantum theory at the same level of clarity as classical physics. In particular, I have argued that the intelligibility of classical physics rests on it possessing a coherent tripartite structure (which spans from an overarching conception of physics and physical reality—the mechano-geometric conception—all the way to specific mathematical theories) rather than it comporting with our experience of the everyday physical world. Accordingly, I contend that quantum theory will achieve a comparable level of intelligibility once it is embedded within an analogous tripartite structure.

In this paper, I have described a reconstruction-based strategy for elucidating quantum theory which, unlike most traditional elucidative methodologies, has the potential of yielding such a tripartite structure. As I have argued, reconstruction is part of the natural life-cycle of physical theories, and is likely to yield the most radical insights when used to investigate theories whose formalism was largely shaped by regularities in the phenomena of interest rather than by *a priori* general principles. As quantum theory is such a theory *par excellence*, its reconstruction is particularly apt. Yet, for the first eighty or so years of its existence, no compelling reconstruction was available. It is thus a blessing that, at this moment in time, we have access to numerous detailed reconstructions of many key parts of the quantum formalism, a bonanza that is largely (although not exclusively) due to the informational perspective on physical theories.

The next step in the elucidative strategy is the *interpretation* of suitably-chosen reconstructions. The *goal* of such interpretation is clear: we wish for a precisely-articulated quantum conceptual framework, analogous to the classical conceptual framework, together with an overarching conception of physical reality which renders intelligible the particular assumptions within that framework.

Although the interpretation of quantum reconstructions is in its early stages, certain insights have already been obtained. In particular, I have sketched some general implications of reconstructions of the abstract quantum formalism (aQF), and have summarized some rather striking implications drawn from a recent reconstruction of the quantum symmetrization algorithm (QSA).

The prospects for further rapid progress in the interpretation of reconstructions is unclear. The fundamental limiting factor is the sheer cognitive difficulty of bringing to bear a reflective or philosophical mindset onto reconstructions that are often articulated in a distinctive mathematical framework, often employ unfamiliar mathematical machinery (such as functional equations or the geometry of convex sets), and posit numerous physical principles whose primary *raison d'être* is their sufficiency for reconstruction rather than their perspicuity. Social factors pose an additional barrier. For example, very few workers in the foundations of physics who develop reconstructions have gone on to interpret them in a deliberate manner. That may, in part, reflect a lack of philosophical training which, in turn, may reflect the scarcity of philosophical education, particularly at the pre-university level. It may also reflect the modern-day gulf between the physics and philosophy communities. In addition, from my outsider's perspective, the philosophical landscape itself appears rather balkanized, and only certain very specific traditions appear to have maintained a strong connection to mainstream physics. Furthermore, in the mainstream philosophy of physics, reconstructive work does not seem to be taken seriously as an alternative pathway to the interpretation of quantum theory. Instead, traditional interpretations (particularly those that seek to preserve as much as possible

of the classical conceptual framework) continue to hold sway.

My own recent efforts to interpret my previous reconstructive work have convinced me that many disparate areas of philosophy—including metaphysics (both analytical and other forms), continental philosophy, and non-western philosophy (such as Mādhyamika philosophy)—have a great deal to offer. For example, Husserl’s phenomenology seems to strongly resonate with the operational reconstructive strategy for elucidation of physical theories: the emphasis on operational procedures reflects the centrality of moment-by-moment perceptual experience (Berghofer and Wiltsche 2019), while the importance of reconstruction echoes Husserl’s genetic phenomenology (Berghofer and Wiltsche 2020, §1.2.8), *viz.* the importance of becoming fully aware of the metaphysical assumptions implicit in the highly-sedimented practices which we unwittingly absorb through inculturation (Berghofer et al. 2021).

In this regard, it is very encouraging to witness ‘metaphysics of physics’ sessions at a recent physics conference, conferences on the metaphysics of science, and the vibrant conference series ‘Phenomenological approaches to physics’ initiated by Harald Wiltsche and Philipp Berghofer. It is my hope that such a confluence of physicists and philosophers from diverse traditions, acting in a concerted manner to interpret the fruits of the quantum reconstruction program, will finally succeed in lifting the quantum veil.

Aczél, J. (1966). *Lectures on Functional Equations and their Application*. Academic Press.

Aczél, J. and Dhombres, J. (1989). *Functional equations in several variables*. Cambridge Univ. Press.

Barrett, J. (2006). Information processing in generalized probabilistic theories. arXiv:quant-ph/0508211v3.

Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics*, 1:195–200.

Berghofer, P., Goyal, P., and Wiltsche, H. A. (2021). Husserl, the mathematization of nature, and the informational reconstruction of quantum theory. *Cont. Phil. Rev.*, 54:413–436.

Berghofer, P. and Wiltsche, H. A. (2019). The co-presentational character of perception. In Limbeck, C. and Stadler, F., editors, *The Philosophy of Perception and Observation*. De Gruyter.

Berghofer, P. and Wiltsche, H. A. (2020). Phenomenological approaches to physics. Mapping the field. In Wiltsche, H. A. and Berghofer, P., editors, *Phenomenological Approaches to Physics*, pages 1–47.

Bohm, D. (1952). A suggested interpretation of the quantum theory in terms of “hidden” variables, I and II. *Phys. Rev.*, 85:166–93.

Bohr, N. (1928). The quantum postulate and the recent development of atomic theory. *Nature*, 121:580–590. Reprinted in Bohr (1961).

Bohr, N. (1961). *Atomic theory and the description of Nature*. Cambridge Univ. Press.

Caulton, A. (2014). Qualitative individuation in permutation-invariant quantum mechanics. <https://arxiv.org/abs/1409.0247>.

Chiribella, G., Perinotti, P., and D’Ariano, G. M. (2011). Informational derivation of quantum theory. *Phys. Rev. A*, 84:012311.

Cramer, J. G. (1986). The transactional interpretation of quantum mechanics. *Rev. Mod. Phys.*, 58:647–688.

Darrigol, O. (2004). The mystery of the Einstein–Poincaré connection. *Isis*, 95:614–626.

Darrigol, O. (2005). The genesis of the theory of relativity. *Séminaire Poincaré*, 1:1–22.

Darrigol, O. (2019). Deducing Newton’s second law from relativity principles: a forgotten history. *Archive for History of Exact Sciences*.

Descartes, R. (1982). *Principles of Philosophy*. Kluwer.

- Everett, H. (1957). Relative state formulation of quantum mechanics. *Rev. Mod. Phys.*, 29(3):452–462.
- Feynman, R. P. (1948). Space-time approach to non-relativistic quantum mechanics. *Rev. Mod. Phys.*, 20:367.
- Fuchs, C. A. (2002). Quantum mechanics as quantum information. [quant-ph/0205039](https://arxiv.org/abs/quant-ph/0205039).
- Fuchs, C. A. (2017). Notwithstanding Bohr, the reasons for QBism. *Mind and Matter*, 15(2):245–300.
- Ghirardi, G., Marinatto, L., and Weber, T. (2002). Entanglement and properties of composite quantum systems: A conceptual and mathematical analysis. *J. Stat. Phys.*, 108(1/2).
- Goyal, P. (2008). Information-geometric reconstruction of quantum theory. *Phys. Rev. A*, 78(5):052120.
- Goyal, P. (2012). Information physics—towards a new conception of physical reality. *Information*, 3(4):567.
- Goyal, P. (2014). Derivation of quantum theory from Feynman’s rules. *Phys. Rev. A*, 89:032120.
- Goyal, P. (2015). Informational approach to the quantum symmetrization postulate. *New J. Phys.*, 17:013043.
- Goyal, P. (2019). Persistence and nonpersistence as complementary models of identical quantum particles. *New J. Phys.*, 21:063031.
- Goyal, P. (2020). Derivation of classical mechanics in an energetic framework via conservation and relativity. *Found. Phys.*, 50:1426–1479.
- Goyal, P. (2022). The nature of identical quantum particles: Insights from an operational reconstruction of the symmetrization postulate. Submitted (*draft available on request*).
- Goyal, P. and Knuth, K. H. (2011). Quantum theory and probability theory: their relationship and origin in symmetry. *Symmetry*, 3:171–206.
- Goyal, P., Knuth, K. H., and Skilling, J. (2010). Origin of complex quantum amplitudes and Feynman’s rules. *Phys. Rev. A*, 81:022109.
- Grinbaum, A. (2007a). Reconstructing instead of interpreting quantum theory. *Philosophy of Science*, 74:761–774.
- Grinbaum, A. (2007b). Reconstruction of quantum theory. *Brit. J. Phil. Sci.*, 58:387–408.
- Hardy, L. (2001a). Quantum theory from five reasonable axioms. <https://arxiv.org/abs/quant-ph/0101012>.
- Hardy, L. (2001b). Why quantum theory? Contribution to NATO Advanced Research Workshop “Modality, Probability, and Bell’s Theorem”, Cracow, Poland, 2001.
- Hardy, L. (2013). Reconstructing quantum theory. <https://arxiv.org/abs/1303.1538>.
- Heelan, P. A. (2016). *The Observable. Heisenberg’s Philosophy of Quantum Mechanics*. Peter Lang.
- Hertz, H. (1899). *The principles of mechanics presented in a new form*. Macmillan & Co. English translation.
- Kochen, S. and Specker, E. P. (1967). *J. Math. and Mech.*, 17:59.
- Lakatos, I. (1978). Falsification and the methodology of scientific research programmes. In Worrall, J. and Currie, G., editors, *The Methodology of Scientific Research Programmes*, pages 8–101. Cambridge Univ. Press.
- Lange, M. (2011). Why do forces add vectorially? A forgotten controversy in the foundations of classical mechanics. *Am. J. Phys.*, 79:380–388.
- Lützen, J. (2005). *Mechanistic Images in Geometric Form: Heinrich Hertz’s ‘Principles of Mechanics’*. OUP.
- Mach, E. (1919). *The Science of Mechanics: A Critical and Historical Account of Its Development*. Open Court Pub. Co., 4th edition.
- Pauli, W. (1998). Exclusion principle and quantum mechanics. In *Nobel lectures in physics (1942–1962)*. World Scientific.
- Schütz, J. R. (1897). Prinzip der absoluten Erhaltung der Energie. *Nachr. v. d. Gesellschaft der Wissenschaften zu Göttingen, Math.-Phy. Kl.*, 2:110–123.
- Schwarz, J. M., Stapp, H. P., and Beauregard, M. (2004). Quantum physics in neuroscience and psychology: a neurophysical model of mind-brain interaction. *Phil. Trans. Roy. Soc. B*.
- Scott, W. L. (1970). *The conflict between atomism and conservation theory 1644 to 1960*. Macdonald (London), Elsevier (NY).

- Wheeler, J. A. (1989). It from bit. In *Proceedings of the 3rd International Symposium on the Foundations of Quantum Mechanics, Tokyo*.
- Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In Zurek, W. H., editor, *Complexity, Entropy, and the Physics of Information*. Addison-Wesley.
- Wigner, E. P. (1960). The unreasonable effectiveness of mathematics in the natural sciences. *Comm. in Pure and Applied Math*, 13:1-14.
- Wootters, W. K. (1980). *The acquisition of information from quantum measurements*. PhD thesis, University of Texas at Austin.