Quantum reality: a pragmaticized neo-Kantian approach[⋆,⋆⋆]

Florian J. Boge[a]

[a]Interdisciplinary Centre for Science and Technology Studies (IZWT), Bergische Universität Wuppertal, Gaußstr. 20, 42119 Wuppertal, Germany

ABSTRACT

Despite remarkable efforts, it remains notoriously difficult to equip quantum theory with a coherent ontology. Hence, Healey (2017, 12) has recently suggested that "quantum theory has no physical ontology and states no facts about physical objects or events", and Fuchs et al. (2014, 752) similarly hold that "quantum mechanics itself does not deal directly with the objective world". While intriguing, these positions either raise the question of how talk of 'physical reality' can even remain meaningful, or they must ultimately embrace a hidden variables-view, in tension with their original project. I here offer a neo-Kantian alternative. In particular, I will show how constitutive elements in the sense of Reichenbach (1920) and Friedman (1999, 2001) can be identified within quantum theory, through considerations of symmetries that allow the constitution of a 'quantum reality', without invoking any notion of a radically mind-independent reality. The resulting conception will inherit elements from pragmatist and 'QBist' approaches, but also differ from them in crucial respects. Furthermore, going beyond the Friedmanian program, I will show how non-fundamental and approximate symmetries can be relevant for identifying constitutive principles.

1. Introduction

1.1. Intriguing developments and their limitations

With its many fruitful applications, quantum theory (QT) has revolutionized both science and the global economy (Kleppner and Jackiw, 2000, 893) but at the same time managed to notoriously escape our intuitive grasp. Of the several well-known proposals for interpreting the formalism in terms of some sort of ontology (e.g. Wallace, 2012b; Dürr et al., 2012; Bassi and Ghirardi, 2003), each has its fundamental difficulties. For instance, "none of the various extant suggestions for *Bohmian quantum field theories* [have delivered,] say, the cross-section for electron-electron scattering, calculated to loop order where renormalisation matters" (Wallace, 2020a, 97; emphasis added). Yet our best predictions in particle physics rely crucially on these very calculations. Similarly, *collapse* interpretations single out *position* as the variable 'collapsed to'; but this is hardly compatible with *decoherence* in the relativistic regime (cf. Wallace, 2012a, 4589), i.e., the gradual vanishing of interference under specific circumstances. Yet essentially all predictions relevant for the Large Hadron Collider rely crucially on the fact that quarks inside protons scattering at high energies are sufficiently decohered in the *momentum* basis (Schwartz, 2014, 674; and below).

The Everett interpretation would be free of such problems if it could coherently recover the Born rule, but that is all but uncontroversial. The approach championed by Deutsch (1999) and Wallace (2012b) has been argued to be threatened by circularity (Baker, 2007), and to suffer from an endorsement of untenable decision-theoretic axioms (Dizadji-Bahmani, 2013; Maudlin, 2014; Boge, 2018). The approaches by Zurek (2005) and Carroll and Sebens (2014, 2018) either rely on *branch-counting*—something incompatible with decoherence, and even with the classical probability calculus (Wallace, 2012b)—, or otherwise, equally invoke dubious rationality principles (Dawid and Friederich, 2019).

These and similar concerns have certainly fueled the development of recent alternatives to interpretations that are straightforwardly 'realist', in the sense of attempting to directly assign referents to elements of the formalism. For instance, according to Healey (2015, 2), there are specific conditions for any quantum state $|\psi\rangle$, its *backing conditions*, under which an agent is warranted to assign it to a system S. The quantum state then offers advice regarding another set of conditions, its *advice conditions* (ibid.). This advice will be offered in terms of the probabilistic information provided by the Born rule, $\Pr(O \in \Delta) = \langle\psi|\Pi_\Delta|\psi\rangle$, where Π_Δ is a projection on $\mathcal{H}$ that corresponds to the range of values Δ for observable O.

In this way, the advice concerns the credibility of the *non-quantum* magnitude claim 'the value of O lies in range Δ'. It does not concern anything peculiar to the quantum formalism. So $|\psi\rangle$ merely functions as an "informational bridge" between backing and advice conditions (Healey, 2015, 2)—it is not a *beable* (Bell, 1976) of the theory, and there are good reasons to suspect that neither are other candidate elements of the formalism, like the local operators $\hat{\phi}(x)$ of quantum field theory (Healey, 2017, 231 ff. and references therein). Rather, "it is the function of *magnitude claims* to repre-

[⋆]This is a post-peer-review, pre-copyedit version of an article accepted for publication in *Studies in History and Philosophy of Science*.

[⋆⋆]©2021. (cc) BY-NC-ND This work is made available under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

fjboge@gmail.com (F. J. Boge)
ORCID(s): 0000-0002-1030-3393 (F. J. Boge)

sent elements of physical reality." (Healey, 2020, 386; emph. added)

In a somewhat similar vein, "[a] QBist[1] takes quantum mechanics to be [...]a very powerful tool that any agent can use to organize her own experience. [...] But quantum mechanics itself does not deal directly with the objective world; it deals with the experiences of that objective world that belong to whatever particular agent is making use of the quantum theory." (Fuchs et al., 2014, 750) And finally Friederich (2015, 8), whose position is inspired by QBism and closely related to Healey's (see also Lewis, 2020), develops "an account of quantum theory [...] that is meant to dissolve [its] problems" in a way that builds on the 'therapeutic' aspects of Wittgenstein's philosophy, and hence cures us of the need to even raise certain interpretive questions.

Now, if Healey, Friederich, and the QBists are right, then so was Bohr; at least in his much quoted claim that "[t]here is no quantum world." (Petersen, 1963, 12) However, if QT thus "has no physical ontology and states no facts about physical objects or events" (Healey, 2017, 12), where else are we supposed to get these from?

QBism focuses on single agents' experiences – which brings it 'dangerously' close to a position that could be identified as solipsistic (Norsen, 2016; Earman, 2019). Furthermore, a "QBist view of science" means that "everything any of us knows about the world is constructed out of his or her individual private experience" (Fuchs et al., 2014, 753). But in QBist writings, we find no *method* that would allow an agent to construct a reality out of her experience. And even less so any basis for inferring the existence of entities, corresponding to certain constructs, like other agents or "'particles' [...] that come to Alice and Bob from a common source S" (Fuchs et al., 2014, 752).

Both Healey and Friederich, on the other hand, lay great emphasis on the aforementioned non-quantum magnitude claims. These are supposed to reflect "statements about entities and magnitudes acknowledged by the rest of physics" (Healey, 2017, 137). But following Wallace (2020b, 386), one may wonder whether either Friederich or

> Healey fully succeeds in saying what the non-quantum physical magnitudes actually are. [...][For instance,] Healey [...] recounts standard cosmology. But it's opaque how this can be recovered in his framework. A 'quantum fluctuation' is a property of the state; on Healey's account, this is just advice to the agent, and can't be causally responsible for anything, least of all features of the early universe billions of years before the agent's birth.

[1] 'QBism' is the now-famous term characterizing the project chiefly undertaken by Fuchs. It originally derives form 'Quantum Bayesianism', but was divorced from that term some time ago.

Furthermore, for Healey (2017, 226), "decoherence is a valuable way to gauge the significance of magnitude claims"; it even provides a "progressive definition of content" to which, however, there is "no natural limit such that one could say that, when this limit is reached, a statement [...] is simply true because one has finally succeeded in establishing a kind of natural language-world correspondence relation in virtue of which the statement correctly represents some radically mind- and language-independent state of affairs." (Healey, 2012, 747)

But many claims which Healey is a realist about, such as those to abstract properties like 'strangeness' or 'color-charge', are not "in any way picked out by decoherence" (Wallace, 2020b, 386). And frankly, if there is no language-world correspondence, it becomes hard to see how one could explicate the semantic utility of QT by reference to, say, "patterns of statistical regularity [...] sufficiently stable to be modeled by Born probabilities" that "*would exist in a world without agents*" (Healey, 2017, 207; emph. added).

It appears, then, that all these positions face a dilemma: Either they are forced to discard the concept of reality as altogether meaningless, reducible to the experiences of a single agent, or something to be cured of; or, on the other horn, to embrace a hidden variables-view.

Neither is willing to face the first horn, and there is some evidence that at least Healey and Friederich face the second. For instance, Friederich (2015, 165) claims that, despite the host of well-known no-go theorems (Leifer, 2014, for an overview), "there is [...] no reason to doubt that appealing and uncontrived assignments of sharp values to all observables are possible". Similarly, Healey (2020, 144) writes:

> If Bohr was right that acceptance of quantum theory requires acknowledgment of the limits this puts on our abilities to speak meaningfully about the physical world, perhaps Einstein was right to hold out the hope that these limits may be transcended as quantum theory is succeeded by an even more successful theory that gives us an approximately true, literal story of what the physical world is like.

And "what else does 'hidden variables' mean" than "factors not found in the theory originally"(van Fraassen, 1991, 243)?

It is obviously true that no no-go theorem exists which rules out all hidden variable theories tout court. But why, if one ultimately opts for hidden variables, buy into a specific pragmatist account of meaning that has the power to render certain claims – and certainly also questions as to their truth – literally *meaningless*? Why develop a therapeutic attitude that should remove the desire to ask these questions? Appealing to the possibility of future theories with further variables as a means for establishing a thorough sort of realism seems to undermine the very motivation for pursuing either Healey's or Friederich's project. Furthermore, one might find it objectionable that this future theory will inevitably inherit certain key features from QT that, for all we know, cannot be reproduced by any model that allows a joint probability distribution over all its dynamical variables (see Jennings and Leifer, 2016).

1.2. Outline of the project

These concerns are enough for me to suggest an alternative route to 'reality in spite of QT', which is in many ways inspired by Heleyanism, Friederichism, and QBism, but also departs from them in crucial ways.[2] Building on the constructivist flavor present in QBism, I will here follow the neo-Kantianism of Friedman (1999, 2001) in taking what is objectively real to be exhausted by relevant symmetries of a given theory. Since 'being', or 'being real', is thereby tied to a specific theory, I do not mean anything radically mind-independent by it. 'Quantum reality', in other words, is nothing but experience successfully synthesized through the application of certain invariant concepts that arise from the quantum formalism.

This means biting the bullet hard: if we give up on our best scientific theories referring to a mind-independent reality in this way, we might have to accept loosing our grip on mind-independent reality altogether; even if we accept that

> our leading scientific theories are the best foundation and starting point we have not only for uncovering new and unexpected phenomena, but also for opening up new areas and paths of inquiry, and in guiding ourselves to the even more powerful conceptions of natural domains that will ultimately replace the ones we now have. (Stanford, 2006, 207)

But I bite this bullet in a way that allows us to attribute a definite meaning to the word 'reality', based directly on features of the quantum formalism, and hence avoids, in my view, the undesirable features of the aforementioned approaches.

In brief, I will argue that certain elements of QT are constitutively a priori, i.e., constitutive of QT's subject matter and a priori relative to its content. These constitutive a priori elements can be found out by considering QT's symmetries, where by 'symmetry' I mean some transformation that leaves some elements (say, states, variables, interactions...) invariant. However, certainly not every symmetry of a theory is constitutive in this sense. Hence, I will identity two types of 'pragmatism' – though nothing as elaborate as embraced by Healey or Friederich – relevant to (a) theory construction and (b) articulation that help us identify the relevant symmetries and what it is that they say.

The structure of the paper is as follows. Sect. 2 first provides a general introduction to this usage of symmetries and its (faint but important) connections to Kant's actual philosophy. Sect. 2.1 discusses the example preferred by both Friedman and Reichenbach (i.e., relativity), to see the construction by means of symmetries at work, which in Sect. 2.2 is extended, in a (well-known) first approach, to QT.

Sect. 3 then identifies the two senses of pragmatism, and gives examples as to how they are employed. Sect. 4 gives deeper consideration to non-fundamental, and merely approximate symmetries which nevertheless function constitutively. This means going an important step beyond the proper Friedmanian program, and conveys the central insights of the paper. It also indicates some commonalities with the positions of Healey and Friederich, as decoherence is seen to play a particularly important role. Finally, in Sect. 5, I will comment on positions other than Friedman's (in particular: various structuralisms), and how they relate to mine.

2. Symmetries and the constitutive a priori: a first approach

Let us begin by recalling the key problem of the original Kantian doctrine. Kant (CPR, A158/B197) was famously concerned with "conditions of the possibility of experience in general" that would "at the same time" be "conditions of the possibility of the objects of experience themselves, and thus possess objective validity in a synthetical judgment *a priori*." He also declared space and time "pure forms of our sensibility" (A494/B522), which allowed him to endow the principles of Euclidean geometry with such a synthetic a priori validity (A47/B64).

With the rise of non-Euclidean geometries in the 19th century, these views became untenable (Friedman, 1999, 6). Reichenbach (1920, 46), however, realized that "the notion of an a priori has two distinct meanings in Kant. Firstly, it means something like 'apodictically valid', 'valid for all times', and secondly it means 'constitutive for the concept of an object'." (my translation—FJB)

Both Reichenbach (1920) and Friedman (1999, 2001) suggest to dispose of the first meaning while keeping the second intact. This would allow us to sort out the 'axioms of coordination' of a given theory Θ that "must be laid down antecedently to ensure [...] empirical well-definedness in the first place." (Friedman, 1999, 61) These contrast with 'axioms of connection' that provide "empirical laws in the usual sense involving terms and concepts that are already sufficiently well defined." (ibid.) The former ones lay out what features of the theory are *constitutively* a priori, i.e., constitutive of the subject matter of Θ and a priori relative to Θ's content. So they are "structurally and functionally [...] that without

[2] The basic ideas have already been developed in Boge (2018, Sect. 7.4). I here significantly extend the treatment, in particular by paying attention to the relevance conditions for symmetries.

which the rest of a theory would lack content." (Howard, 2010, 337)

In addition, Friedman (1999, 66) suggests that what is constitutively a priori can be found out by determining the theory's invariants under a relevant transformation group. Equivalently, we may speak of finding a theory's *symmetries*, where "the symmetry of a 'something' (a figure, an equation,...) is defined in terms of its invariance with respect to a specified transformation group, its symmetry group." (Castellani, 2003, 322)

The group-theoretic notion of symmetries may be too narrow. For instance, Guay and Hepburn (2009) explore groupoids, i.e. (essentially), connected collections of groups, as the appropriate means for exhibiting symmetry-properties of, e.g., the Rydberg spectrum. I will here, in fact, not pay close attention to the fundamental mathematical description of symmetry transformations, but rather consider any transformation of a theory Θ a 'symmetry' that leaves some of its elements (states, variables, interactions...) invariant.

That symmetries or invariants are also an essential part of the original Kantian methodology can be gathered from the writings of various Kant scholars. Schrader (1951, 520), for instance, claims that: "One of Kant's proofs of the a priority of space and time rests upon the fact that space and time are universal and invariant in a way that specific sensed objects are not." Similarly, Allison (2015, 340) has it that for Kant, "transcendental consciousness[...] grounds objectivity by providing a standpoint that is invariant with respect to the contingent contents of the various episodes of empirical consciousness." At the same time, this introduces an element into Kant's philosophy that by today's standards might be considered *anti-realist*: "On Kant's account, [...] *we* provide [...] a *general* form—space, time, and the categories—for which sensible intuition supplies the only content that this invariant humanly possible 'scheme' can possibly have." (Rosenberg, 2005, 250; orig. emph.)

This Kantian use of symmetries has been identified as a precursor to, or version of, what phenomenologists call 'eidetic variation' (e.g. Wiesing, 2014, 64). But since "[t]he starting point for an eidetic variation is *conscious reflection on one's own mental state*" (ibid., 65; emph. added), the connection between Kant's use of symmetries and that in modern physics may seem contrived.

I believe that such a verdict would be too quick: The main difference is that Kant was reflecting on our *manifest* image, trying to establish its objective foundations, which were then also supposed to generalize to objective foundations of science. This latter project turned out to be untenable, however, and physicists today are interested rather in a specifically *scientific* image; one that can accommodate a huge range of observations still unavailable to Kant. Defining this image with the aid of symmetries may require relativizing constitutive efforts to a given theory (or a set thereof), and invoking 'exotic', mathematically formulated symmetries. But this does not make for a fundamental difference in the general approach: specify a method of variation, consider what remains constant under all allowed variations, and build an objective empirical reality (an image) based on these constants. That symmetries can be regarded in this broadly Kantian way will be a firm presupposition for what follows.

2.1. The classic(al) case

To see constitution by means of symmetries at work, take the following example. In general relativity (GR), spacetime is represented by an ordered tuple $(\mathcal{M}, T^{(i)})_{i\in I}$. Here $\mathcal{M}$ is a (Lorentzian) manifold, a topological space that can be covered by subsets $U \subseteq \mathcal{M}$ which can be mapped continuously into $\mathbb{R}^n$, and the $T^{(i)}$ are tensor fields defined thereon. In GR, these are the metric tensor g_{ab} and the stress-energy tensor ψ_{ab}, specifying local geometric properties and the distribution of stress-energy and matter, respectively.

Reichenbach (1920), as recounted by Friedman (1999, 66; emph. omitted), had in essence concluded that "only the underlying topology and manifold structure remain constitutively a priori", while "physical geometry (the metric of physical space) is no longer constitutive." However, it is not clear whether, or in what sense, this verdict can hold. Consider two ordered tuples $(\mathcal{M}, g_{ab}, \psi_{ab})$, $(\mathcal{N}, \tilde{g}_{ab}, \tilde{\psi}_{ab})$, and assume that $\mathcal{M}$ and $\mathcal{N}$ are related by a *diffeomorphism* ϕ, a bijective C^∞-map with ϕ^{-1} also C^∞, where $\mathcal{N}$ is the image of $\mathcal{M}$ under ϕ. In this case, the two tuples are taken to represent the same spacetime, because if one is a solution to the Einstein field equations, so is the other.

Since diffeomorphisms correspond to smooth deformations, $\mathcal{M}$ and $\mathcal{N}$ are topologically identical. Diffeomorphisms also preserve the manifolds' (maximal) smooth structures, so we have some right to identifying $\mathcal{M}$ and $\mathcal{N}$ as 'one and the same object' (cf. Nakahara, 2003, 180). In contrast, the tensor fields, including the metric tensor g_{ab}, will remain invariant under ϕ only in special cases—for g_{ab}: when ϕ happens to be an isometry.

This may explain the Reichanbachian intuition in modern terminology. However, ϕ will also induce a map ϕ_* (the 'pushforward') that relates $\tilde{T}^{(i)} = \phi_* T^{(i)}$, and the effect of smoothly deforming $\mathcal{M}$ into $\mathcal{N}$ is compensated by the change $T^{(i)} \mapsto \phi_* T^{(i)}$; figuratively, ϕ_* can 'reorganize' the $T^{(i)}$ as needed and so *effect* an isometry between $(\mathcal{M}, g_{ab}, \psi_{ab})$ and $(\mathcal{N}, \tilde{g}_{ab}, \tilde{\psi}_{ab})$. Thus, it is not clear that Reichenbach's verdict on the constitutive aprioricity of the manifold holds, as we obtain the same right to considering $T^{(i)}$ and $\phi_* T^{(i)}$ one and the same object.[3]

[3] I owe special thanks to an anonymous referee for helping me get this point straight. Cf. also Weatherall (2018) for an excellent recent disentanglement of these issues; and cf. there and Weatherall (2016) for reasons to forgo a discussion of (principal) fiber bundles in this connection.

Understandably, the exact message of diffeomorphism invariance in GR is subject to considerable debate. However, what I take from this (all too) brief discussion is that the invariance of the entire *structure* $(\mathcal{M}, g_{ab}, \psi_{ab})$ points us, not to a failure of the general analysis, but to different constitutive elements than Reichenbach had originally though: Spacetime is a set of (dynamic) *geometric relations* between material objects, regardless of the frame in which we chose to represent them *and* any more specific relation between the 'place-times' at which they occur (the choice of manifold within the equivalence class).

There is a basic intuition underlying such symmetry-based analyses that may be parsed as follows: As we saw, the elements of the symmetry group represent transformations of some 'object' (e.g. a space-time continuum), and if that object has features which occur when 'viewed from all angles or perspectives' as provided by the group's elements, then these features must *really* pertain to the object, not just as a 'perspectival effect'.

"Stated so bluntly," however, this description verges on what van Fraassen (2006, 292) coined a "sophistry [that] will not take in anyone." Hence, note that, in contrast to the position actually attacked by van Fraassen – ontic structural realism (short: OSR; see Ladyman and Ross, 2007) –, objects, in particular those of scientific discourse, are here intimately tied to a conceptual scheme, so we can infer "nothing whatsoever about the things in themselves that may ground them." (CPR, A49/B66) Furthermore, with van Fraassen (2006, ibid.), we also concede that "relevance is contextual." It will hence be the burden of the paper to investigate the relevance of certain symmetries of the quantum formalism for the question of reality constitution—albeit not in full completion, but in part rather by (verbal) ostension.

2.2. Towards a quantum constitution

In a first approach, then, the above considerations may be transferred to QT as follows. As is well known, (closed) linear subspaces of a system S's Hilbert space $\mathcal{H}_S$ can be understood as a representation of S's properties. Given an observable O with associated self-adjoint operator
$\hat{O} = \int_{\sigma(O)} o\Pi(\mathrm{d}o)$ ($\sigma(O)$ O's spectrum), the subspace of vectors that correspond to a range $\Delta \subseteq \sigma(O)$ may represent the proposition that O takes on values in Δ on S. As first pointed out by Birkhoff and von Neumann (1936), intersection, closed linear span, and orthocomplementation may then be used to define algebraic relations that could be interpreted in terms of the logical operations $\vee, \wedge, \neg$.

Now the symmetry group for an elementary system S in non-relativistic QT is the Galilei group, $\mathcal{G}$, which has an irreducible projective unitary representation $U_{\mathcal{G}}$ on $\mathcal{H}_S$, parametrized by particle mass m and spin quantum number s (cf. Lévy-Leblond, 1967). However, the fact that $|\psi_S\rangle$ cannot be a simultaneous eigenstate of all projections of $\mathcal{H}_S$, i.e., that not all the observables that can be defined on S are *comeasurable*, implies that the algebra of propositions so generated is not *Boolean*.

Since only the entire lattice is invariant under $U_{\mathcal{G}}$, an elementary quantum system is that whichever *carries* the properties from the non-Boolean lattice, "in all situations which can be obtained by Galilei-transformations." (Mittelstaedt, 2009, 857) But since a Boolean sublattice can only be defined by considering a set of *commuting* self-adjoint operators, at any given time this only allows us to attribute to it "a class of mutually commensurable properties" (ibid.).

Essentially, this is a highbrow restatement of a sense or aspect of Bohrian complementarity. However, it allows us to say that what it is to *be* an elementary system in QT: an entity that allows the context-sensitive[4] *ascription* of properties—*not* something with a pre-defined set of properties evolving over time.

The treatment neatly generalizes to special relativistic QT if one considers the Poincaré group, $\mathcal{P}$, instead of $\mathcal{G}$; a consequence of a famous paper by Wigner (1939). A "slight generalization" (Emch and Piron, 1963, 470) then again allows talk of subspaces representing properties, the total algebraic structure of which is invariant under a unitary representation of $\mathcal{P}$.

Moreover, since spin and mass are "invariant under relativistic transformations[...] the kinematic characteristics of *types* of free particles can be obtained from spatio-temporal invariance." (Auyang, 1995, 37; emph. added) As Streater (1988, 144; orig. emph.) puts it:

> Wigner [...] did not merely say that a particle is well described by such a representation:[...] a particle *is* a pair $(\mathcal{H}, U_{[m,s]})$, where $\mathcal{H}$ is a Hilbert space and U is a unitary continuous action of $\mathcal{P}$ on $\mathcal{H}$, obeying $U(a,\Lambda)$ $U(b,M) = \omega U(a+\Lambda b, \Lambda M)$, where $a, b \in \mathbb{R}^4$ are space-time vectors, and Λ, M are Lorentz matrices, and where $[m, s]$ are the mass and spin.

From a stronger realist point of view, this statement may be puzzling. But while this might not be Streater's intended reading, it becomes perfectly comprehensible against the Kantian backdrop developed here: what it is to *be* an electron (say),[5] according to QT, is to have

[4] I do not mean anything specifically physical by 'context' here. Rather, I mean specific (experimental or observational) conditions under which a given agent would be prone to associate some state vector or density matrix to the system under study (as is evidently the case in physical practice), so as to be able to associate a subset of properties consistent with the state-assignment to it (like, say, its having a certain well-defined momentum).

[5] An anonymous referee has rightly pointed out to me that Wigner's is insufficient as a more fine-grained classification that also distinguishes e.g. electrons from positrons. It is hence vital to note that this classification is not exhaustive, and that further invariants, such as the conserved charge that can be identified, on account of Noether's first theorem, in the Lagrangian of quantum electrodynamics by means of (global) U(1) invariance, are necessary for that sake.

a specific (rest) mass and kind of spin, and to allow a number of assertions relative to a given context (according to the Boolean sub-lattices contextually implied by $\mathcal{H}$)—not a tiny charged 'corpuscle' traveling on a trajectory in space, nor even an emergent feature of the global, branching wave-field that fills the universe.

3. Neo-Kantianism pragmaticized

Our first approach has brought us some way towards the constitution of a quantum reality, but it ultimately has limited appeal. Subspaces of a Hilbert space (or projections onto these) play a subordinate role in special-relativistic QT, whereas field operators $\hat{\phi}(x)$ take center stage (e.g. Streater, 1988, 137–41, for a recap of the historical and systematic reasons). In quantum field theory (QFT), moreover, the relevant (Fock) spaces are not generally *unitarily equivalent*,[6] meaning that neither states from, nor operators on, two spaces $\mathcal{F}, \mathcal{F}'$ are generally connected by a unitary map U. Hence, QFT implies limitations for successfully constituting objects by means of unitarily preserved algebras of properties.

The above constitutive efforts, moreover, concern only *free* systems, which have applications only under highly idealizing conditions: *interactions* (e.g. between measuring apparatus and system studied) will be exploited in any conceivable experiment. Algebraic approaches to QFT add to these worries: Haag's theorem effectively trivializes QFTs in which free particles exist, in the sense that the field operators in such a QFT cannot enter into *any* interactions.

All these problems, however, relate to the fact that the above treatment is too 'purist': It only says something about *fundamental* concepts of the theory, independently of applications. This is reflected, for instance, in the fact that the problems raised by Haag's theorem can be avoided by considering only *asymptotic* freedom, i.e., creation operators for particles that can be treated as free *for all practical purposes* (Bain, 2000, for an excellent exposition).

I suggest that *pragmatic* considerations are required for an extension of the Kantian analysis to the theory as actually employed. Before delving into the details, let us consider in what ways pragmatic considerations may enter.

3.1. Pragmatism at a foundational level

Note first that Bohr himself has been called a pragmaticized Kantian by Folse (1994, 121–2):

> Pragmatized Kantians defend their claims to knowledge through appeal to the pragmatic virtues of the categories under which the content of experience is subsumed. [...] Bohr's work in philosophy is in effect simply this: a campaign to revise the limits of application of key concepts in the physicist's synthesis of the experiences which form the empirical basis of our knowledge of the atomic domain.

The upshot is that pragmatic considerations may figure in the very *formation* of new, radically different theories, whenever accepted ones face (severe) failure: We revise our conceptual inventory in part according to given *practical needs*, not conceptual and overall empirical plausibility. Let us call this 'pragmatism$_a$'. It can be seen as closely connected to the present Kantian agenda as follows.

Friedman (2001, Chap. 3) raises the question of how the conceptual change associated with the acceptance of a new, radically different theory can be 'communicatively rational' in the sense that it still allows us "to engage in argumentative deliberation or reasoning with one another aimed at bringing about an agreement or consensus of opinion." (Friedman, 2001, 54; emph. added) His response is that

> the new constitutive framework is a quite deliberate modification or transformation of the old constitutive framework, *developed against the backdrop of a common set of problems, conceptualizations, and concerns*.

Compare this to Bohr's (1928, 580) allusions to an "irrationality" involved in the postulation of a quantum of action, while simultaneously considering QT, with its operators satisfying many of the relations of corresponding classical variables, a "rational generalisation of the classical theories" (ibid., 584). The same problem set, enriched by new problems that render an old approach insufficient, may be tackled by a rather deliberate modification of the constitutive framework. Nevertheless, this may still count as rational insofar as one preserves the benefits of the old framework and connects to the new one in a sufficiently communicable way.

A connection of pragmatism$_a$ to the suggested analysis in terms of symmetries can be established as follows. Consider, as an example, how we ended up with local gauge invariance in QFT. At the classical level, the electrodynamic Hamiltonian has to include the electromagnetic vector potential A_μ in order to ensure that the phenomenologically valid Lorentz-Coulomb force law $\boldsymbol{F}_{EM} = e(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B})$ derives from it in the usual way. But A_μ is only unique up to a gauge transformation $A_\mu \mapsto A_\mu + \partial_\mu \chi$, with χ essentially arbitrary.

As was first realized by Fock (1926), when the Hamiltonian is promoted to an operator in QT, the shift in A_μ requires a simultaneous transformation $\psi \mapsto \psi e^{-\imath \chi e}$ (in natural units) for the entire (Schrödinger- or Dirac-) equation to remain invariant, where χ now acts as a

I do not see this as a principled problem, though, as I will find it necessary anyways to go beyond the 'interaction-free limit' (in which Wigner's results are derived).

[6] Cf. however Ruetsche (2011, 57 ff.) for a discussion of unitary inequivalence already in *non-relativistic* QT, as can occur when the configuration space of the system has a topology different from $\mathbb{R}^n$.

local phase (Jackson and Okun, 2001, for the historical details).

But why do we even use a Hamiltonian formulation in the first place, rather than confining attention to forces and measurable quantities such as ***E*** and ***B***?

> The sum of kinetic and potential energy was originally introduced into mechanics as an auxiliary, purely mathematical entity, arising as a first integral of the Newtonian equations of motion for systems subject to conservative forces. But as a result of the formulation of the general principle of the conservation of energy and its incorporation in the science of thermodynamics (the First Law) it came to be regarded as possessing ontological significance in its own right. (Redhead, 2003, 128–9)

Thus, entirely (*a*-)pragmatic considerations, such as a maintenance of the Lorentz-Coulomb force law in what was at first simply a convenient formalism, lie at the root of local U(1)-symmetry in quantum electrodynamics.

Shortly after Fock's initial contribution, local gauge invariance was, however, "declared a general principle and 'consecrated' by Hermann Weyl[...]." (Jackson and Okun, 2001, 663) As such, it has served as a remarkably productive prescription in the formation of new QFTs. Hence, a new constitutive principle is created on initially pragmatic grounds:

> there is ultimately *no compelling logic* for the vital leap to a local phase invariance [...]. Nevertheless, the gauge principle – deriving interactions from the requirement of local phase invariance – provides a satisfying conceptual unification of the interactions present in the Standard Model. (Aitchison and Hey, 2012, 61; emph. added)

The *intuitive* content of local gauge invariance is by no means as straightforwardly extracted as with the symmetries discussed above (also Lyre, 2009, 215). However, it is evident that local gauge transformations only play a role regarding the *interactions* between fields: Fields can be gauged in such a way that local characteristics like the local phase of scalar fields or local amplitudes of vector fields are arbitrary, *so long as these are appropriately correlated*. This sheds light at least on its *target*: what it means for 'fields' (local magnitudes) to *interact* (to contribute to measurable values in concert), according to gauge invariant QFTs like the Standard Model (SM),[7] is to do so in a way that only depends on the relation between these characteristics at any spacetime point. I take it to be a corollary of this assessment that the terms of the Lagrangian form the 'ontological bedrock' of the theory, in the sense of that which is supposed to be objectively real on its account—*not* the gaugable fields with (all of) their (partly gauge-dependent) individual properties. Hence, it is certain relations entered into by the field operators that offer us the primary clue as to what we should take reality according to QT to be, not the operators so related themselves.

The main similarity with diffeomorphism invariance in GR lies in the fact that this relegates many features of the formal entities involved to the domain of the perspectival and non-objective: Just as choosing particular coordinates means subscribing to a particular viewpoint of the situation that must be carefully taken into account in one's ultimate result does working in a particular gauge mean taking a particular view on the given field operator(s) that must be just as carefully taken into account.

3.2. Pragmatism at the within-theory level

Pragmatism$_a$ is but one sort of pragmatism relevant for us here. The other sort, 'pragmatism$_b$', is closely connected to Ruetsche's (2011, 147; emph. altered) 'unpristine' approach to theory-*interpretation*:

> The doctrine of unpristine interpretation allows that the contingent application of theories does not *merely select* among some preconfigured set of their contents, but *genuinely alters* their contents. It follows that there can be an *a posteriori*, even a pragmatic, dimension to content specification, and that physical possibility is not monolithic but kaleidoscopic.

Ruetsche's subsequent discussion is entangled with questions of modality, which I do not intend to address. Moreover, the linguistic dimension of pragmatism inherent in Healey's favoured version obviously comes to mind again. But the pragmatism suggested here is neither as radical nor directly (and exclusively) tied to inferential practices.[8]

What I intend pragmatism$_b$ to stand for is merely that the conditions of application have a say in the specification of the very *content* of a theory. More concretely, this means that they may figure in *singling out* objective features – thereby possibly bestowing relevance on certain symmetries – of an existing, 'finalized' theory.

As an example, consider the almost elusive notion of a 'point particle' in QFT. As is well known, there are several results (Malament, 1996; Halvorson and Clifton,

[7] Recall that this is the union of electroweak theory and quantum chromodynamics.

[8] In fact, Healey's endorsement of Brandom's (1994, 2000) particular brand of pragmatism is one major reason for me not to follow down the Healeyan route exactly: As Fodor and Lepore (2001, 2007) argue in detail, there does not seem to be any way to recover compositionality within Brandom's position; the fairly uncontroversial feature that we can understand an indefinite number of sentences in virtue of our understanding of their constituents. To my knowledge, Brandom has not even responded to the 2007 paper anywhere in print.

2002) that suggest that a rigorous notion of a point particle as some sort of entity localizable in a narrow volume of space is ill-defined. Yet physicists still like to employ this notion without further ado. As explication of the concept, one sometimes finds reference to the fact that interactions are modeled in QFT by coupling field operators at spacetime *points* (Hatfield, 1992; Cao, 1999). This, however, really only establishes "the pointlike nature of the *particle interactions*: we construct interaction Hamiltonians by multiplying the relevant fields at exactly the same spacetime point." (Duncan, 2012, 164; orig. emph.) It is "*not* a statement about our ability to localize the physical characteristics [...] at a dimensionless spatial point" (ibid.).

A similar but more specific point is made by Falkenburg (2007) and (less explicitly) Muller (2014). In scattering scenarios, the central quantity to be evaluated is the (differential) *cross section* $\frac{d\sigma}{d\Omega}$, which provides a prediction for the particles to be expected in a solid angle $d\Omega$. In QFT, cross sections include a matrix element that may be expressed in terms of so called 'form factors'. These form factors, in turn, generally depend on the momentum transfer in the scattering.

The 'Rosenbluth' cross section for elastic electron-proton scattering, for instance, has two form factors $F_{1/2}(q^2)$ (q the momentum transfer between two scattering particles):

$$\frac{d\sigma}{d\Omega} \propto \frac{\cos^2(\frac{\theta}{2})}{\sin^6(\frac{\theta}{2})}\left[F_1^2 - \frac{q^2}{4M^2}[2(F_1+2MF_2)^2\tan^2\left(\frac{\theta}{2}\right)+(2MF_2)^2]\right] \tag{1}$$

Such a cross section can be made dimensionless by multiplying through by an appropriate quantity that has the physical dimension 1/area. If a dimensionless cross section then turns out to be *scaling invariant*, i.e. "does *not* depend on any length, one concludes that the scattering center and the probe particles are structureless or *pointlike*." (Falkenburg, 2007, 133; orig. emph.) The intuition being that *structured* particles are 'breakable into finer pieces' through 'harder smashing'. Hence, "a particle has [...] structure – i.e. is *not* a point particle – if and only if the functions $F_1(q^2)$ and/or $F_2(q^2)$ are *not* constant." (Drell and Zachariasen, 1961, 8; emph. added) Contrapositively, it *is* 'pointlike' just in case they *are* constants, and the particle's scattering behavior does not depend on the momentum it receives.

Setting $F_1 = 1$ and $F_2 = 0$ in the Rosenbluth formula (Eq. 1), for instance, we retrieve the scattering cross section for electron-quark scattering, when the quark is assumed free and not part of the proton's structure (Schwartz, 2014, 674). Hence, the extent to which this cross section agrees with experiment justifies the assumption of essentially free,[9] 'pointlike' quarks inside the proton.

How can we even *talk* about 'colliding' particles, though, when pointlikeness has nothing to with them being confined to trajectories? That is a headache we need to defer to Sect. 4.3. The lesson to be drawn here is this: what it is to be a 'point particle', according to the SM as applied to scattering scenarios, is to be an entity associated with a scaling-invariant cross section. Or in other words: one whose scattering behavior remains the same when 'smashed together', no matter how hard, with something equally 'pointlike', so that harder smashing never reveals any substructure.

Clearly, there is a symmetry involved here but that symmetry is certainly nothing like the fundamental symmetries of the SM (Poincaré invariance and local SU(3) ×SU(2)×U(1) gauge-invariance). Nevertheless, paying attention to a symmetry that only figures (or even emerges) in application can apparently lead to the identification of a constitutive principle. What drives physicists in considering scaling invariance in this constitutive capacity? It can only be the desire to coherently employ the theory in a way that facilitates talk of 'point particles', even though, on the face of it, the formalism seems to treat only of continuous, extended entities ('fields').[10]

4. Non-fundamental and approximate symmetries

In the preceding section, I gave an example of a symmetry that was singled out as having a constitutive function by a context of application: Applying what could be naïvely construed as a model of interacting fields to a situation in which it is more helpful to talk of scattering particles, some of which are 'pointlike' in nature, we encountered a symmetry that allowed us to say what 'pointlikeness' really means, and in what sense we should assume 'pointlike particles' to exist. Despite its constitutive function, this is, however, arguably, not a fundamental symmetry of the SM, and traditionally, only fundamental symmetries have been considered to possess constitutive relevance (see Sect. 2.1).

Now, 'fundamental' is sometimes used as opposed to 'accidental', but it would be certainly too restrictive to think of scaling invariance as something that "can be changed or even abandoned without any collateral changes in nature, since nothing depends on [it]" (Kosso, 2000, 119), or that it is a "dynamical accident[...] having no fundamental physical significance." (Redhead, 1975, 81) Hence, what, if neither 'accidental' nor obviously fundamental, should we take this symmetry to be?

[9] This asymptotic freedom for proximal quarks is a prerequisite for establishing their point-likeness; for otherwise any probe would interact with the whole (composite) hadron and no scaling invariance should be expected in the first place. Note also that the relevant form factors for the electron-proton scattering used to test this prediction do depend on the fraction x of the proton's momentum that each quark carries. The scaling invariance is then retrieved as the approximate independence from q^2 *at a given* x.

[10] It should be obvious that the notorious measurement problem does not allow any straightforward, literal construal of particles as 'excitations of the fields'; see also Halvorson and Clifton (2002, 207).

Furthermore, I noted that scaling invariance only holds approximately for electron-quark scattering, since quarks are never perfectly free. Hence, in some cases, even approximate symmetries can apparently have a constitutive function. In fact, there is one particularly important approximate symmetry, namely that involved in quantum decoherence, which allows us to talk as if a classical constitution of objects, according to which they travel on fixed trajectories at well-defined speeds, was often still appropriate. I will later offer a more detailed account of how decoherence can function as a bridging principle in this way, but I will now first turn to the relevant notion of non-fundamentality.

4.1. Constitutive non-fundamental symmetries as symmetries of sub-theories

What kind of a symmetry is scaling invariance if it is neither fundamental nor accidental? In a first approach, consider the meaning of the word 'theory'. Following a relatively non-committal version of the semantic view, we can think of theories at bottom as families of models that are constrained by a number of general laws (e.g. Giere, 1988). Accordingly, the SM might be seen as a theory of the sub-atomic domain, which prescribes the general laws of interaction and the dynamics of quantized fields. Any detailed model of, say, a particular kind of interaction between specific fields, subject to particular boundary conditions, would be constrained by these general laws.

Now consider, in contrast, the notion of a 'perturbation theory'. At bottom, perturbation theory tells us in which ways an interaction-free process is altered if small contributions of an interaction (small 'perturbations') are taken into account. However, there is not just one perturbation theory: In non-relativistic QT, perturbation series are just as important as in QFT, but non-relativistic QT neither needs to pay attention to Poincaré invariance nor appeal to local operators. Furthermore, even within the SM, there are different possible ways of doing a perturbation series: One can either use the Lagrangian as defined in terms of the bare parameters (masses, charges), or renormalize them first. The renormalized coupling will then depend on some scale, and this conveys automatic justification on the use a given perturbation series, in the sense of defining where the series is 'valid' to the extent that higher order terms will contribute less and less. This justification needs to be somehow recaptured when one starts off with the bare Lagrangian (Schwartz, 2014, 341).

Neither of these latter two perturbation theories operates at the same level as the SM, but they are still more general than any concrete model formulated within the SM. I hence suggest to consider such theories as *sub-theories* of a more encompassing one: They stipulate further principles that are relevant for defining specific models that are otherwise not sufficiently constrained. Renormalized and un-renormalized perturbation theory should thus be viewed as two rival sub-theories of the SM, initially specifying different constraints on what models of a given perturbed process should look like. Non-relativistic QT, on the other hand, contains a whole other perturbation theory that only 'saves' (and extends) the laws of non-relativistic QT.

I believe that this notion of a sub-theory can be fruitfully employed to understand the possibility of symmetries that possess constitutive relevance but are in a clear sense neither fundamental nor accidental. In particular, I suggest to view scaling invariance as a constitutive principle *of the scattering theory that is a sub-theory of the SM*. It is thus not fundamental *for the SM itself*. But it still tells us that, what it means for quarks to be pointlike, is to exhibit a scaling invariant behavior *in scattering experiments*.

This explains the sense in which scaling invariance can function constitutively, namely as a constitutive principle of a sub-theory of the SM. And it also tells us in what sense it is non-fundamental; namely that it is not fundamental for the SM itself. But this still does not explain in what sense it is 'non-accidental'. To gain a clearer understanding of this, consider Fletcher's (2019, 18) recent account of accidental symmetries, which aims at combining the virtues of several other accounts without buying into their vices. According to Fletcher, a symmetry T (approximate or exact) is accidental in some theory Θ, relative to a theory Θ', when Θ' is (i) more encompassing in the sense of having (a) possibilities in its scope that are in a relevant sense similar to those of Θ, and that (b) also explain these, but (ii) T is not a symmetry of Θ'. The idea is, in essence, that Θ' 'saves the phenomena' of Θ, but can at the same time explain Θ's apparent symmetry *T away*; say, as a "dynamical accident[...] having no fundamental physical significance." (Redhead, 1975, 81)

Now it should be obvious that we cannot interpret Θ' as the SM and Θ as the scattering theory contained in it, for the possibilities of the SM's scattering theory are wholly preserved in the SM—not explained away by relevantly similar possibilities. I hence suggest to think of scaling invariance not as accidental, but rather as *local*; not in the sense of being a function of spacetime points, but as being a symmetry over a subspace of the possibilities of the SM, as defined by the principles of one of its sub-theories.

This implies an interesting spin on the relation between pragmatism$_a$ and pragmatism$_b$: Applying the SM to scattering scenarios can lead to new, local constitutive principles, induced by the needs connected to the given application (e.g., the need of defining additional principles in applying the SM to scattering-scenarios). Hence, pragmatism$_b$, insofar as it bestows relevance upon certain symmetries not fundamental for the more encompassing theory, opens the way for a local pragmatism$_a$: it aids in figuring out the objective features of sub-theories of a more encompassing theory (like the SM), once that theory has itself been established on partly *a*-pragmatic grounds.[11]

4.2. Decoherence as a bridging principle

Now I also noted that scaling invariance is not just non-fundamental but, in the case of quarks, also approximate. There is a reason from within the SM why we should expect this to be so (the confinement of quarks into hadrons). But insofar as we can at least conceptually (though not mathematically) extend the theory to conditions under which we could regard quarks as free, which also allows us to identify conditions (very, though not arbitrarily, high energies) under which we can treat them as 'free for all practical purposes', we can equivalently treat this symmetry as exact for these very purposes, and then harvest the content provided by it as a (local) constitutive principle.

An even more fruitful application of this treatment of approximate symmetry lies in decoherence theory. Decoherence, to recall, essentially amounts to the vanishing of interference terms and the selection of a preferred basis in virtue of interactions between systems. It is often regarded as the 'emergence of classicality' due to the influence of a system's environment (e.g. Joos et al., 2003; Schlosshauer, 2007), even though not all bases picked out by models predicting decoherence will correspond to a kind of classical description.

Consider a system S that couples to an environment E, equipped with quantum states $|\mathcal{E}_j\rangle$. For modeling the dynamics of both, it is often justified to consider only the interaction Hamiltonian (cf. Schlosshauer, 2007, 77). After suitable interaction, one would then end up with a state $|\Psi_{SE}\rangle = \sum_j \alpha_j |S_j\rangle |\mathcal{E}_j\rangle$, and the projector $|\Psi_{SE}\rangle\langle\Psi_{SE}|$ onto this state defines a pure state density matrix $\hat{\rho}_{SE}$. If we compute the effective density matrix for S by tracing out E, we obtain:

$$\begin{aligned}\hat{\rho}_S = \mathrm{Tr}_E(\hat{\rho}_{SE}) &= \sum_{i,j} \alpha_i \alpha_j^* |S_i\rangle\langle S_j| \,\mathrm{Tr}\big(|\mathcal{E}_i\rangle\langle\mathcal{E}_j|\big) \\ = \ldots &= \sum_{i,j} \alpha_i \alpha_j^* |S_i\rangle\langle S_j| \,\langle\mathcal{E}_j|\mathcal{E}_i\rangle\,. \end{aligned} \tag{2}$$

Assume also that $\langle\mathcal{E}_j|\mathcal{E}_i\rangle \approx 0$ for $i \neq j$. Then we get $\hat{\rho}_S \approx \sum_j |\alpha_j|^2 |S_j\rangle\langle S_j|$, which looks like a statistical mixture in which the 'true' quantum state could be simply unknown.

The assumption of orthogonal environment states for different system states is not merely academic: In multiple-scattering models (cf. Joos et al. 2003, 64 ff.; Schlosshauer 2007, 119 ff.), one retrieves, in the limit of large times, a differential equation

$$\frac{\partial}{\partial t}\rho_S(\boldsymbol{x},\boldsymbol{x}';t) = -F(\boldsymbol{x}-\boldsymbol{x}')\rho_S(0) \tag{3}$$

in the position representation, which is solved by an exponential in the *decoherence factor* $F(\boldsymbol{x}-\boldsymbol{x}')$, scaled by the initial density matrix. The effect is that the evolution thus prescribed exponentially damps the off-diagonal terms of the density matrix. However, the derivation involves an intermediate time scale that is required to be much larger than the time of an individual scatter (but almost zero as compared to time scales of human interest). This, in turn, has the effect that the damping will never be perfect, and that the Fourier-transform will be of a similar form (with little interference between individual momentum states).

Note the various elements of the *b*-pragmatic type here: It is relative to time scales that are almost zero as compared to time scales of *human interest* that one can derive the result. Furthermore, the interference will be small, even negligible, relative to *humanly achievable* precision, but strictly speaking not zero. However, when all this is ignored, decoherence induced by scattering can usually be taken to have the effect of driving quantum states into states that are resemble mixtures of position eigenstates, sufficiently 'unsharp' so as to also allow for the simultaneous ascription of an 'unsharp' momentum. The resulting state hence closely resembles a statistical mixture of particles with definite, but only approximately known, trajectories.

How does all this relate to considerations of invariance? In fact, the observable $\hat{O}$ preferred by decoherence, as well as its associated eigenbasis, are selected by $\hat{O}$'s approximate invariance under the continuing influence of the environment, as expressed by $[\hat{O},\hat{H}_{\mathrm{int}}] \approx 0$, with $\hat{H}_{\mathrm{int}}$ the Hamiltonian that models the interaction between S and E (cf. Zurek, 1982, 1869). As soon as the approximation can be considered 'valid', we may constitute an object on the basis of the preferred observables.

There are various ways of spelling out the 'validity' of such an approximate invariance. In the present context, this notion is best captured exactly by the requirement that, relative to achievable precision and accuracy, there be no *measurable* difference between any decohered state $\hat{\rho}_d$ and a corresponding mixture $\hat{\rho}_m$ of eigenstates $|o_j\rangle\langle o_j|$ of $\hat{O}$. *In principle*, these could always be distinguished in the long run (cf. d'Espagnat, 1990, 1154 ff.), though in some (if not most) cases only in a universe whose material content, so far as we know, vastly exceeds that of our own one (cf. Omnès, 1994, 307–9).

To make this a little more precise, recall that $[\hat{O},\hat{H}_{\mathrm{int}}] \approx 0$ really means that for any reference state $|\psi\rangle \in \mathcal{H}_{SE}$ it holds that $(\hat{O}\otimes\mathbb{I})\hat{H}_{\mathrm{int}}|\psi\rangle \approx \hat{H}_{\mathrm{int}}(\hat{O}\otimes\mathbb{I})|\psi\rangle$ ($\mathbb{I}$ the identity on the part of the joint Hilbert space that doesn't concern S). This has the particular consequence that eigenstates of the interaction will align with eigenstates of $\hat{O}\otimes\mathbb{I}$. So given that $\hat{H}_{\mathrm{int}}$ is the infinitesimal generator of the transformation $\hat{U}_{\mathrm{int}}(\delta) = 1 - (\imath/\hbar)\hat{H}_{\mathrm{int}}\delta + \ldots$, the prescription tells us that, over time, the eigenstates of $\hat{H}_{\mathrm{int}}$ will be transformed into eigenstates of $\hat{O}\otimes\mathbb{I}$ (if only at infinity).

[11] I owe thanks to anonymous referee for the correct expression of this relation between both pragmatisms.

Now following, again, Fletcher's recent, detailed account of approximate symmetries,[12] we can interpret this more precisely as follows. For Fletcher (2019, 4), the physical possibility-space $\mathcal{S}$ of a system, i.e., "the possible ways that states of affairs could be arranged for the system", is endowed with a set $\mathcal{R}$ of similarity (i.e., reflexive, symmetric) relations $\sim\subseteq \mathcal{S}^2$. If these are induced by a pseudometric $d : \mathcal{S}^2 \to [0, \infty)$, then one can (partially) order them in strength according to the value $\epsilon \in [0, \infty)$ that d assigns. More generally, if $\sim'\subseteq\sim$, then $\sim'$ may be said to be at least as (or, in case of proper inclusion: more) *discriminating*, as it allows fewer (or at most equally many) physical possibilities to be similar.

However, if a transformation $T : \mathcal{S} \to \mathcal{S}$ has the property that, for a given $\sim$, $T(s) \sim s$ for all $s \in \mathcal{S}$, then T may be a considered an $\sim$-*approximate symmetry* (it transforms states considered similar into one another). Moreover, if for given $\sim\in \mathcal{R}$, $T'(s) \sim s$ for all $s \in \mathcal{S}$, but there is a s' such that $T'(s) \nsim s'$, then T' may equally be said to be more discriminating (it doesn't preserve all the similarities).

The point is that, for δ large enough, there will be some ϵ representing the aforementioned experimental conditions such that $\|\hat{\rho}_d - \hat{\rho}_m\|_{\text{tr}}/2 = \text{tr}[|\hat{\rho}_d - \hat{\rho}_m|]/2 \leq \epsilon$, where δ paramterizes progressing time, $\hat{\rho}_d$ is the reduced system-density matrix that evolves according to $\hat{U}_{\text{int}}(\delta)$, $\hat{\rho}_m$ is a mixture of eigenstates of $\hat{O}$, and $|\hat{X}| = (\hat{X}^\dagger\hat{X})^{1/2}$ ($\|\hat{X} - \hat{Y}\|_{\text{tr}}/2$ the distance measure induced by the trace-norm).

Given that the nature of scatterings is actually discrete, we can take this to mean that, over time, the symmetry of the quantum state changes in a succession $T, T', T'', \ldots$, where, up to a certain point, the T become more and more discriminating against states that differ substantially from (mixtures of) eigenstates of the 'pointer observable' $\hat{O}$. Or in yet other words: Successively, the (changing) approximate symmetry induced by the interaction will only preserve states that lie within ϵ of the trace distance of a certain mixture $\hat{\rho}_m$.

However, as was pointed out above, if $\hat{O}$ represents position, this will usually be true (for some $\epsilon' \approx \epsilon$) also of a mixture of *momentum* eigenstates $|\boldsymbol{p}\rangle\langle\boldsymbol{p}|$. For instance, assume that the initial state of S is a superposition of two Gaussian wave packets in (one-dimensional) position-space – i.e., the classic toy model for a realistic superposition of being 'here and there at once' – and that the decoherence factor $F(\boldsymbol{x} - \boldsymbol{x}')$ can be given as $\Lambda(\boldsymbol{x} - \boldsymbol{x}')^2$ (Λ a characteristic damping-rate). It can then be shown (cf. Joos et al., 2003, 93 ff., and references therein) that the Wigner transform

$$W(x, p) = \frac{1}{\pi\hbar}\int_{-\infty}^{\infty} \mathrm{d}y \exp\left(\frac{2\imath py}{\hbar}\right)\rho(x - y, x + y) \qquad (4)$$

becomes strictly positive. Moreover, by definition, a Wigner transform is normalized to unity over all phase-space points (x, p). Hence, under these conditions, it 'mimics' the properties of a classical phase space density.

Well in line with Fletcher's (2019, 12) verdict that "the pertinent empirical equivalence-preserving symmetries of the old theory can be explained as being merely approximately empirical equivalence-preserving in the reducing theory", we can take it that the approximate symmetry thus involved in decoherence induces a *bridging principle*: The positivity of a phase-space density corresponding to (4) is invariant under time-translations in classical statistical mechanics, but this is merely approximately correct for the corresponding Wigner transform of the density matrix.

However, for all practical purposes, and relative to a specified set of physical conditions, we thus *dynamically* retrieve something akin to a phase-space distribution, i.e., something that can be treated as specifying the statistical properties of an ensemble of particles following definite trajectories at definite speeds. Hence, under these conditions, QT allows us to constitute (effectively) 'classical particles'.

The b-pragmatic elements to this were already sorted out above: We usually need several considerations of what can be neglected *in practice* in even deriving the relevant formal results, and we always need to include a cutoff at which we consider the symmetry to be close enough to exact so that we can talk as if the system in question 'really was' classical. In other words: Scientists' frequent talk of, e.g., elementary particles in ways that seems strictly incompatible with QT is sanctioned by the fact that there are – sometimes quite heuristic, but usually quantitative – arguments that allow us to *switch constitutive systems*; i.e., to *talk as if* the old, essentially abandoned constitutive system was still appropriate

This assessment of decoherence and classical properties should once more be compared to a spacetime example: in Euclidean geometry triangles' angles sum to π, and this may be considered a constitutive feature insofar as it reflects a scaling invariance of the sides: any simultaneous rescaling that does not mess up the triangle will not change the angular sum. In *non-Euclidean* spaces, however, this is not generally true.

Thus, consider a possible world in which physical space behaves geometrically like the surface of a solid sphere. Angles would here sum to $\pi + A/a^2$ instead, with A the triangle's area and a the sphere's radius (e.g. Hartle, 2003, p. 18). Conscious beings in that possible world could be unaware of this if all scales they encountered were small compared to the size of the sphere.

[12] I have some reservations about Fletcher's account, beyond those expressed by himself (cf. his Fn. 13). For instance, realistic talk of 'models of a theory' may presuppose a notion of similarity among models (Giere, 1988), whence it is unclear to me whether the account is not threatened by circularity. However, I am not aware of any better worked-out account of approximate symmetries, so I shall suppress these worries in this paper.

Once aware of the non-Euclideanness of their space, however, they could use the *approximate* invariance of the sum at sufficiently small scales to retrospectively justify their original Euclidean image, and to retain it for many practical purposes.

There is an obvious similarity to decoherence, but the disconnect between quantum and classical reality is far more radical: the curved space can be embedded in the Euclidean one and is somewhat imaginable from within the Euclidean framework. But what is left, from the point of view of classical physics and everyday life thinking, if we strip away such properties as *location* and *speed* from physical objects?

4.3. Colliding particles

Despite these counter-intuitive properties, it is exactly decoherence that also aids in making sense of 'colliding' particles; something that we had left undetermined above. At the Large Hadron Collider (LHC), the initial particles are protons, and a standard locution is that these are prepared in 'bunches' by a magnet system, crossing every 25 ns. Without giving a numerical estimate, we can follow the arguments in Joos et al. (2003, Sect. 4.1, and references therein), and assume that the interaction with the magnetic fields alone will exert a sufficiently decohering influence to justify the assertion of more or less localized bunches traveling at more or less well-defined momenta.

The same certainly holds for all particles whose trajectories are reconstructed from the signals received from the detector: the manifold interactions with the detector (GEANT Collaboration, 2016, for an impressive overview) will drive stable products from the goings on in the 'beam pipe' into quasi-classical states. But this really only means that, relative to magnet system and detector, decoherence-style arguments allow us to constitute the relevant particles as quasi-classical entities crossing at rather well-defined spacetime points with well-defined energy-momenta, or whose passage one may reconstruct in terms of trajectories and (transverse) momenta respectively.

On the level of the interacting *partons* (quarks and gluons), however, relativistic four-momenta stick out as relevant quantities, as they essentially provide all the information required to identify different event types. It is a general assumption, proven only for a small range of cases (cf. Collins et al. 1989, 1–2; Schwartz 2014, 685), that a pp cross section with final state particles c and d factorizes as

$$\sigma(pp \to cd) = \sum_{a,b} \int_0^1 \int_0^1 \mathrm{d}x_a \, \mathrm{d}x_b \, f_a(x_a; Q^2) f_b(x_b; Q^2) \sigma(ab \to cd). \quad (5)$$

The sum ranges over parton flavors, $\sigma(ab \to cd)$ are parton-level cross sections, computed perturbatively in quantum chromodynamics (QCD), and the $f_j(x_j; Q^2)$ are *parton density functions* (pdfs), which, under certain circumstances, can be interpreted as a probability density for finding a parton of flavor j at fraction x_j of the proton's momentum (up to small transverse fluctuations; cf. Schwartz, 2014, 696–7).

Such factorizations only hold up to some power of Λ/Q, where Λ is the scale below which perturbative QCD calculations break down, and Q "some characteristic high-energy scale in the process." (Schwartz, 2014, 685) In this context, one often reads that partons 'interact incoherently', and Schwartz (2014, 674; emph. altered) indeed claims that: "To actually prove that [...] *decoherence occurs* amounts to a proof of factorization." Informally, this can be seen from the fact that the pdfs contain information on the surrounding remainder of the proton (a parton's 'environment'), and that the interference between different partonic momentum states is suppressed as $\Lambda/Q \to 0$, under which condition also the probabilistic interpretation of the pdfs becomes possible. However, the elementary cross sections contain (mod-squared) field-theoretical matrix elements, so when $\Lambda/Q \approx 0$, the computation roughly corresponds to one in which the initial state is a mixed state over all possibly contributing parton flavors and momenta.

The lesson to be learned here is that, in a context such as the LHC, where much of our evidence regarding the sub-nuclear constitution of matter comes from, and from which QT gains much of its support as being the empirically most successful theory to date, decoherence-style arguments not only allow us to treat protons as more or less classical, and so facilitate the vast amount of loose talk about proton bunches crossing—which talk, of course, is crucial for the purposes of experimental design and evaluation. Decoherence arguments *also* facilitate a treatment of protons as statistical ensembles of partons in well-defined momentum-eigenstates; i.e., 'plane wave' states which are maximally indefinite regarding spacetime information. This, in turn, is crucial for the computation of matrix elements on the parton level in the experimentally preferred basis, and so ultimately for the ability to make sense of reconstructed particle tracks in the detector in terms of the presence and properties of elementary particles like the Higgs.

When combined with various pragmatic considerations that allow one to tickle out certain quantitatively precise statements, the quantum formalism itself thus provides the conditions of its own *use*: It determines under which conditions we can switch constitutive systems and constitute protons 'classically', but also which *non-classical* aspects we should pay attention to (partons scattering in well-defined, essentially free four-momentum states, not proton-structure), and to what extent (weight assigned by the pdfs).

However, notice a hitch to this assessment: It presupposes that the quantum state is viewed as inherently *probabilistic*, for we do not retrieve *one single* state that corresponds to a particle on some trajectory (Bell, 1990). Hence, only if we accept, along the lines of Healey (2015) or Fuchs et al. (2014), that QT's main function is to provide advice about expected future experiences or the applicability of certain claims about physical magnitudes can we move from a quantum to a classical reality constitution.

4.4. The objectivity of quantum correlations

We here finally turn to an assessment of quantum entanglement, "*the* characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought." (Schrödinger, 1935, 555; orig. emph) A prominent example of an entangled state is the spin singlet $|\chi\rangle = 2^{-1/2}(|\uparrow\rangle|\downarrow\rangle - |\downarrow\rangle|\uparrow\rangle)$, recently used in a(n acclaimed) loophole-free violation of Bell-inequalities with spinful systems separated by 1.3 kilometers (see Hensen et al., 2015).

As is well known, $|\chi\rangle$ predicts a perfect anti-correlation between spin values whenever measured along the same axis. Any unitary transformation U that does not couple to the spin-degrees of freedom will leave this correlation untouched, which is why it is possible to perform measurements that confirm the predicted statistics: U can model the 'traveling' of two particles in a state $\psi(\boldsymbol{x}_1, \boldsymbol{x}_2)|\chi\rangle$ from a common source towards (say) two Stern-Gerlach magnets, without affecting the $|\chi\rangle$-part. Moreover, the (squared) norms of quantum states are invariant under unitaries ($|U|\chi\rangle|^2 = ||\chi\rangle|^2$), so the perfect anti-correlation $p(\uparrow_{\text{one side}} \mid \downarrow_{\text{other side}}) = 1$ comes out as an objective feature of the particle pair on our Kantian analysis.

However, as is equally well known, no direction of measurement is specified by $|\chi\rangle$; the state is *rotation invariant*. One may see the entire 'mystery' of the quantum correlations as rooted in this fact: regardless of the axis of measurement, we find a perfect match between spin up and down when measured along the same axis. But since there is no way that the two spins can 'arrange' their alignment without violating the causal constraints set up by both relativity *and* standard principles usually assumed in probabilistic theories of causation (cf. Wood and Spekkens, 2015; Näger, 2016), one is forced to assume either non-standard and superluminal or backwards causation (Maudlin, 2011; Näger, 2016; Evans, 2018) or that the correlations themselves are in some sense *uncaused* (Wüthrich, 2014; Gebharter and Retzlaff, 2020).

Ironically,[13] the neo-Kantian analysis favors the latter option: The correlation is an objective feature of reality, but the values of individual spins (or any additional properties λ) that could exert a causal influence on the final values are not. In a qualified sense, we can thus agree with Mermin (1998, 753; emph. omit.): "Correlations have physical reality; that which they correlate does not."

[13] Kant famously attributed a synthetic a priori status to causal closure (e.g. CPR, B134).

Now Portmann and Wüthrich (2007, 849) point out that certain "approaches to quantum gravity[...] suggest that tiny violations of Lorentz group invariance are to be expected. Seen as an implication of rotation invariance, $[p(\uparrow_{\text{one side}} \mid \downarrow_{\text{other side}}) = 1]$[...] would not be warranted any more." For example, when two entangled particles are emitted from a point O in a spacetime curved by the presence of a gravitational field, and their correlation at points $A_{1/2}$ is considered, the stronger-than-classical "correlation may actually arise [...], but only approximately, if the particles are localized around the points O, A_1 , A_2 in regions which are narrow in comparison with the distances OA_1 and OA_2." (von Borzeszkowski and Mensky, 2000, 199) Moreover, "[t]he longer the propagation and the stronger the gravitational field, the poorer is the correlation." (ibid., 202) The reason being that a global notion of spatial direction is ill-defined in a curved spacetime, and that a correspondence between local coordinate frames must be established using parallel transport, particularly along the curve A_1OA_2 (ibid., 198–9).

The perfect (anti-)correlation may, in other words, be spoiled *on theoretical grounds*, on account of a future, more encompassing theory. The degree of divergence from a perfect correlation is rooted in the *non-invariance* of the z-coor-dinate under parallel transport, or the loss of rotation invariance respectively. According to the neo-Kantian analysis, the non-objectivity of global reference frames in a curved spacetime hence surprisingly implies the *non*-objectivity of perfect spin-correlations.

Nevertheless, under most circumstances, these correlations would still be *stronger than classical*—and objectively so, given that their objectivity (in contrast to their *perfectness*) follows from the unitary invariance of the norm, not from rotation invariance of the state. Given the considerations of Sect. 4.3, we may hence invoke the *approximate* rotation invariance in approximately flat spacetimes as bridging between the constitution of stronger-than-classical correlations in a special- and general relativistic quantum reality: According to a special relativistic constitution, perfect quantum correlations are objectively real; according to a general relativistic one, they are only so in the total absence of a gravitational field or over infinitesimal trajectories; but in any case, there will be objective, stronger-than-classical correlations.

5. Connections to other positions in the philosophy of science

Given the importance of symmetries and relations recognized throughout this paper, it should be obvious that there are close affinities to various kinds of *structuralism*. However, as was pointed out already in Sect. 2.1, embracing a Kantian attitude towards 'objects' makes for a clear distinction to OSR. In the words of French (2014, 99-100):

> in following the neo-Kantian in her rejection of objects, the structuralist need not go all the way and follow her down to what she sees as the ultimate ground of objectivity. Instead, the structuralist can resituate objectivity in the laws and principles of our best theories, rather than the putative objects, and the structural realist can take the former as representing features of a mind-independent reality, on the basis of the standard realist arguments (such as the No Miracles Argument).

In contrast, I have here embraced a view (or variant) of Kant's philosophy according to which it is "more properly seen as epistemological or perhaps 'metaepistemological' than as metaphysical in nature" (Allison, 2004, 4), in the sense of informing us about the very conditions of scientific knowledge, not about the structure of the world.

On the face of it, this still leaves room for a close connection to *epistemic* structural realism (short: ESR; see Worrall, 1989), the position that we do have relational knowledge of objects situated in an external reality, but not of their intrinsic properties or natures.

It is true: some Kant-interpreters (Langton, 1998) display Kant's philosophy in essence as a version of ESR. However, given that objectivity was tied here to a constitutive framework and so even remarkable (cor)relations, definitive of the very *subject-matter* of science at one point, may lose their objective character (Sect. 4.4), this is not an appropriate assessment of the position defended here. Rather, the two readings of the implications of Kant's philosophy mostly have in common the acknowledgement of an induced *epistemic modesty*, albeit cashed out in rather different terms (Allison, 2004).

Closer affinities certainly exist to Massimi's neo-Kantian structuralism, which adopts an "internalist perspective about physical reality as dependent on the particular experimental and theoretical circumstances we can avail ourselves of" (Massimi, 2011, 13). However, according to Massimi (2005, 24), there is a "constitutive/regulative dichotomy [that] marks a gulf among Kantian scholars as to whether it is the faculty of understanding with its constitutive principles, or rather the faculty of reason with its regulative principles that is ultimately responsible for the law-governedness of nature."

The constitutive/regulative distinction in Kant's philosophy is extremely subtle (e.g. Friedman, 1992) and cannot be discussed in any serious detail here. However, it should be clear that the invariants isolated in this paper must be associated with a constitutive status which is relative to the given theory. Hence, I do not follow Massimi's exact version of a dynamical Kantianism (which has its roots in Buchdahl and Cassirer), according to which changing scientific principles are merely regulative (see Massimi, 2005, 24).

Finally, I acknowledged an affinity to van Fraassen, in allowing decidedly pragmatic, contextual elements as having a say in the determination of symmetries' relevance. This was, in fact, a crucial move in construing, say, scaling invariance as constitutive of point-particles. Hence, recall the two basic pillars of what van Fraassen (2006, 2010) calls *empricist* structuralism (ES):

> I. Science represents the empirical phenomena as embeddable in certain abstract structures (theoretical models).
>
> II. Those abstract structures are describable only up to structural isomorphism. (van Fraassen, 2010, 238)

I. is van Fraassen's go on the connection between mathematically formulated science and empirical evidence. An embedding is strictly speaking a relation between mathematical structures, meaning that the embedded structure is (or is isomorphic to) a substructure of a larger one.

This raises the question how phenomena could possibly be 'embedded' in mathematical structures. van Fraassen (2010, 240 ff.) is at pains to give an answer to this, for both observed and unobserved phenomena, which answer consists in gesturing at the relation between phenomena and models in concrete examples; in the case of actual observations via a detour through data models and surface models (smoothings of data models). But this defense does not really exceed the verdict of da Costa and French (2003) in their partial structures-approach to truth, embraced also in Bueno's (1999; 2010) version of ES: that, ultimately, "the nature of this relationship lies beyond linguistic expression." (da Costa and French, 2003, 17)

Point II., on the other hand, makes for the decisively 'structuralist' part, as isomorphism leaves behind only 'bare structures'. Formulated thus, however, ES is almost indistinguishable from ES*R*: Worrall (2020, 200; orig. emph.) has recently defended ESR against the infamous Newman-objection by pointing out that "the real Ramsey sentence *turns only theoretical and not observational predicates* into variables and then existentially quantifies over them." This is something easily acceptable for (and in close parallel to) van Fraassen's line of defense.

On the other hand, van Fraassen (2006, 295) notices an

> air of schizophrenia [to ESR that] comes from the fact that all the support given for the [...] claim [that ... scientific knowledge is cumulative in some important respect] is explicitly and admittedly concerned only with an accumulation of empirical knowledge.

The dividing line really is that, in contrast to ESR, ES "is a view not of what nature is like but of what science is." (van Fraassen, 2010, 239) In that respect, the present Kantian approch is indeed closer to ES than to ESR, as it locates on the same side of the metaphysics / epistemology divide. What I take issue with in ES, however, is the point stressed by Nagel (2000, 346), and embraced also by Ladyman (2000, 2010):

> To make the kind of epistemic use of experience that empiricism demands, we need at least the capacity to sort out its deliverances from other products of the mind [...] and this sorting task is [...] a rational enterprise [...] that demands substantive a priori knowledge for its execution.

For, to recall, "Thoughts without content are empty, intuitions [Anschauungen] without concepts [Begriffe] are blind." (CPR, A51/B75)

6. Conclusions

I have offered a road to quantum reality that builds on the suggestion that QT has no referential function, without relying on the problematic features of related projects (Fuchs et al., 2014; Friederich, 2015; Healey, 2017). In a way, my approach reduces the often acclaimed 'strange' character of the quantum formalism, by relating its specific epistemological problems to those of other scientific theories, or even to traditional philosophical discourse.

This was accomplished here by isolating elements of the formalism that may count as having a (relativized) a priori status and then demonstrating how an objective reality can be constructed on account of them. However, the solution may be considered as radical as the problem: For "a priori sources of cognition determine their own boundaries by that very fact (that they are merely conditions of sensibility), namely that they apply to objects only so far as they are considered as appearances, but do not present things in themselves." (CPR, A39/B56)

A subtlety associated with the broadly constructivist nature of the Kantian approach is that, once the rigidity of the Kantian categories is removed, the approach threatens to collapse into a thorough relativism. If correct, this would render past transitions undergone by science expressions of arbitrariness, subjectivity, and a want of rationally comprehensible methodology. This issue is addressed in detail also by Friedman (2001), who counters that the changes between constitutive frameworks may count as rational insofar as they allow for a "*retrospective* communicative rationality", meaning that "practitioners at a later stage are always in a position to understand and rationally to justify—at least in their own terms—all the results of earlier stages." (ibid. 96; orig. emph.)

As we have seen, it is largely decoherence that allows for this retrospective communicative rationality in QT, as it facilitates a sensible bridging between quantum and classical-probabilistic treatments under relevant conditions.

But, one must not overinterpret this facilitation in terms of stronger realist commitments: since "the later constitutive framework employs essentially different constitutive principles", we cannot recover "the classical constitutive framework as such, but only an empirical counterpart to this classical framework formulated within an entirely different constitutive framework." (Friedman, 2001, 98) Hence, given the radicality of the breach between constitutive systems with the advent of QT, the latter may indeed urge of us "a radical revision of our attitude towards the problem of physical reality." (Bohr, 1935, 697)

Acknowledgements Some research for this paper was done during my employment with the research unit The Epistemology of the Large Hadron Collider, funded by the German Research Foundation (DFG; grant FOR 2063). I thank David Glick, Chris Fuchs, Dean Rickles, Radin Dardashti, Richard Healey and two anonymous referees for helpful comments and Robert Harlander for checking the physics parts.

References

Aitchison, I. and Hey, A. (2012). *Gauge Theories in Particle Physics: A Practical Introduction. Volume I*. Taylor & Francis, fourth edition.

Allison, H. E. (2004). *Kant's Transcendental Idealism*. New Haven, London: Yale University Press, revised and enlarged edition.

Allison, H. E. (2015). *Kant's Transcendental Deduction: An Analytical-Historical Commentary*. Oxford, New York: Oxford University Press.

Auyang, S. Y. (1995). *How Is Quantum Field Theory Possible?* Oxford, New York: Oxford University Press.

Bain, J. (2000). Against particle/field duality: Asymptotic particle states and interpolating fields in interacting QFT (or: Who's afraid of Haag's theorem?). *Erkenntnis*, 53(3):375–406.

Baker, D. J. (2007). Measurement outcomes and probability in Everettian quantum mechanics. *Studies In History and Philosophy of Science Part B*, 38(1):153–169.

Bassi, A. and Ghirardi, G. (2003). Dynamical reduction models. *Physics Reports*, 379(5):257–426.

Bell, J. S. (1987[1976]). The theory of local beables. In Bell, J. S., editor, *Speakable and unspeakable in quantum mechanics*, pages 52–62. Cambridge, New York: Cambridge University Press.

Bell, J. S. (1990). "Against 'measurement'". In Miller, A. I., editor, *Sixty-Two Years of Uncertainty. Historical, Philosophical, and Physical Inquiries into the Foundations of Quantum Mechanics*, pages 17–32. New York: Plenum Press.

Birkhoff, G. and von Neumann, J. (1936). The logic of quantum mechanics. *Annals of Mathematics*, 37(4):823–843.

Boge, F. J. (2018). *Quantum Mechanics Between Ontology and Epistemology*. Cham: Springer International Publishing.

Bohr, N. (1928). The quantum postulate and the recent development of atomic theory. *Nature*, 121(3050):580–590.

Bohr, N. (1935). Can quantum-mechanical description of physical reality be considered complete? *Physical review*, 48(8):696–702.

Brandom, R. B. (1994). *Making It Explicit. Reasoning, Representing,* *and Discursive Commitment*. Cambridge, Massachusetts: Harvard University Press.

Brandom, R. B. (2000). *Articulating Reasons. An Introduction to Inferentialism*. Cambridge, Massachusetts: Harvard University Press.

Bueno, O. (1999). What is structural empiricism? Scientific change in an empiricist setting. *Erkenntnis*, 50(1):55–81.

Bueno, O. (2010). Structural empiricism, again. In Bokulich, A. and Bokulich, P., editors, *Scientific structuralism*, pages 81–103. Springer.

Cao, T. Y. (1999). Introduction: Conceptual issues in quantum field theory. In Cao, T. Y., editor, *Conceptual foundations of quantum field theory*, pages 1–27. Cambridge, New York: Cambridge University Press.

Carroll, S. M. and Sebens, C. T. (2014). Many worlds, the Born rule, and self-locating uncertainty. In Struppa, D. C. and Tollaksen, J. M., editors, *Quantum Theory: A Two-Time Success Story*, pages 157–169. Springer.

Carroll, S. M. and Sebens, C. T. (2018). Self-locating uncertainty and the origin of probability in Everettian quantum mechanics. *The British Journal for the Philosophy of Science*, 69(1):25–74.

Castellani, E. (2003). On the meaning of symmetry breaking. In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pages 321–334. Cambridge, New York: Cambridge University Press.

Collins, J. C., Soper, D. E., and Sterman, G. (1989). Factorization of hard processes in QCD. In Mueller, A. H., editor, *Perturbative QCD*, pages 1–91. World Scientific Publishing Co.

da Costa, N. and French, S. (2003). *Science and Partial Truth: A Unitary Approach to Models and Scientific Reasoning*. Oxford, New York: Oxford University Press.

Dawid, R. and Friederich, S. (2019). Epistemic separability and Everettian branches: A critique of Sebens and Carroll. *The British Journal for the Philosophy of Science*. https://doi.org/10.1086/712718.

d'Espagnat, B. (1990). Towards a separable 'empirical reality'? *Foundations of Physics*, 20(10):1147–1172.

Deutsch, D. (1999). Quantum theory of probability and decisions. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 455(1988):3129–3137.

Dizadji-Bahmani, F. (2013). The probability problem in Everettian quantum mechanics persists. *The British Journal for the Philosophy of Science*, 66(2):257–283.

Drell, S. and Zachariasen, F. (1961). *Electromagnetic Structure of Nucleons*. London, Glasgow: Oxford University Press.

Duncan, A. (2012). *The Conceptual Framework of Quantum Field Theory*. Oxford University Press.

Dürr, D., Goldstein, S., and Zanghì, N. (2012). *Quantum Physics Without Quantum Philosophy*. Berlin, Heidelberg: Springer.

Earman, J. (2019). Quantum Bayesianism assessed. *The Monist*, 102(4):403–423.

Emch, G. and Piron, C. (1963). Symmetry in quantum theory. *Journal of Mathematical Physics*, 4(4):469–473.

Evans, P. W. (2018). Quantum causal models, faithfulness, and retrocausality. *The British Journal for the Philosophy of Science*, 69(3):745–774.

Falkenburg, B. (2007). *Particle Metaphysics. A Critical Account of Subatomic Reality*. Berlin, Heidelberg: Springer.

Fletcher, S. C. (2019). An invitation to approximate symmetry, with three applications to intertheoretic relations. *Synthese*, pages 1–21.

Fock, V. (1926). Über die invariante Form der Wellen- und der Bewegungsgleichungen für einen geladenen Massenpunkt. *Zeitschrift für Physik*, 39(2):226–232.

Fodor, J. and Lepore, E. (2001). Brandom's burdens: Compositionality and inferentialism. *Philosophy and Phenomenological Research*, 63(2):465–481.

Fodor, J. and Lepore, E. (2007). Brandom beleaguered. *Philosophy and Phenomenological Research*, 74(3):677–691.

Folse, H. J. (1994). Bohr's framework of complementarity and the realism debate. In Faye, J. and Folse, H. J., editors, *Niels Bohr and Contemporary Philosophy*, pages 119–139. Dordrecht: Springer Science + Business Media.

French, S. (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press.

Friederich, S. (2015). *Interpreting Quantum Theory: A Therapeutic Approach*. Basingstoke: Palgrave Macmillan.

Friedman, M. (1992). Regulative and constitutive. *The Southern Journal of Philosophy*, 30(Supplement):73–102.

Friedman, M. (1999). *Reconsidering Logical Positivism*. Cambridge, New York: Cambridge University Press.

Friedman, M. (2001). *Dynamics of Reason*. Stanford: CSLI Publications.

Fuchs, C. A., Mermin, N. D., and Schack, R. (2014). An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics*, 82(8):749–754.

GEANT Collaboration (2016). Physics reference manual. *GEANT 4*, Release 10.4. https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsReferenceManual/BackupVersions/V10.4/fo/PhysicsReferenceManual.pdf.

Gebharter, A. and Retzlaff, N. (2020). A new proposal how to handle counterexamples to Markov causation à la Cartwright, or: Fixing the chemical factory. *Synthese*, 197(4):1467–1486.

Giere, R. N. (1988). *Explaining science: A cognitive approach*. University of Chicago Press.

Guay, A. and Hepburn, B. (2009). Symmetry and its formalisms: Mathematical aspects. *Philosophy of Science*, 76(2):160–178.

Halvorson, H. and Clifton, R. (2002). No place for particles in relativistic quantum theories? In Kuhlmann, M., Lyre, H., and Wayne, A., editors, *Ontological Aspects of Quantum Field Theory*, pages 181–214. New Jersey, London: World Scientific Publishing.

Hartle, J. B. (2003). *Gravity: An Introduction to Einstein's General Relativity*. San Francisco, Boston: Addison Wesley.

Hatfield, B. (1992). *Quantum Field Theory Of Point Particles And Strings*. Cambridge, Massachusetts: Perseus Books.

Healey, R. (2012). Quantum theory: a pragmatist approach. *The British Journal for the Philosophy of Science*, 63(4):729–771.

Healey, R. (2015). Quantum states as objective informational bridges. *Foundations of Physics*, 47(2):1–13.

Healey, R. (2017). *The Quantum Revolution in Philosophy*. Oxford University Press.

Healey, R. (2020). Pragmatist quantum realism. In French, S. and Saatsi, J., editors, *Scientific Realism and the Quantum*, pages 123–146. Oxford University Press.

Hensen, B., Bernien, H., Dréau, A., Reiserer, A., Kalb, N., Blok, M., Ruitenberg, J., Vermeulen, R., Schouten, R., Abellán, C., et al. (2015). Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575):682–686.

Howard, D. (2010). 'Let me briefly indicate why I do not find this standpoint natural.' Einstein, General Relativity, and the contingent a priori. In Domski, M. and Dickson, M., editors, *Discourse on a New Method, or a Manifesto for a Synthetic Approach to History and Philosophy of Science*, pages 333–356. Chicago, La Salle, Illinois: Open Court.

Jackson, J. D. and Okun, L. B. (2001). Historical roots of gauge invariance. *Reviews of Modern Physics*, 73(3):663.

Jennings, D. and Leifer, M. (2016). No return to classical reality. *Contemporary Physics*, 57(1):60–82.

Joos, E., Zeh, H., Kiefer, C., Giulini, D., Kupsch, J., and Stamatescu, I.-O. (2003). *Decoherence and the Appearance of a Classical World in Quantum Theory*. Berlin, Heidelberg: Springer, second edition.

Kant, I. (1998). *Critique of Pure Reason*. London: MacMillan and Co. Translated by Paul Guyer and Allen W. Wood.

Kleppner, D. and Jackiw, R. (2000). One hundred years of quantum physics. *Science*, 289(5481):893–898.

Kosso, P. (2000). Fundamental and accidental symmetries. *International Studies in the Philosophy of Science*, 14(2):109–121.

Ladyman, J. (2000). What's really wrong with constructive empiricism? Van Fraassen and the metaphysics of modality. *The British Journal for the Philosophy of Science*, 51(4):837–856.

Ladyman, J. (2010). The epistemology of constructive empiricism. In Dicken, P., editor, *Constructive empiricism: epistemology and the philosophy of science*, pages 46–61. Springer.

Ladyman, J. and Ross, D. (2007). *Every Thing Must Go*. Oxford, New York: Oxford University Press.

Langton, R. (1998). *Kantian Humility: Our Ignorance of Things in Themselves*. Clarendon Press.

Leifer, M. (2014). Is the quantum state real? An extended review of ψ-ontology theorems. *Quanta*, 3(1):67–155.

Lévy-Leblond, J.-M. (1967). Nonrelativistic particles and wave equations. *Communications in Mathematical Physics*, 6(4):286–311.

Lewis, P. J. (2020). Quantum mechanics and its (dis)contents. In French, S. and Saatsi, J., editors, *Scientific Realism and the Quantum*, pages 168–184. Oxford University Press.

Lyre, H. (2009). Gauge symmetry. In Greenberger, D., Hentschel, K., and Weinertn, F., editors, *Compendium of Quantum Physics. Concepts, Experiments, History and Philosophy*, pages 248–255. Berlin, Heidelberg: Springer.

Malament, D. (1996). In defense of dogma: Why there cannot be a relativistic quantum mechanics of (localizable) particles. In Clifton, R., editor, *Perspectives on Quantum Reality*, pages 1–10. Dordrecht: Kluwer.

Massimi, M. (2005). *Pauli's Exclusion Principle: The Origin and Validation of a Scientific Principle*. Cambridge University Press.

Massimi, M. (2011). Structural realism: A Neo-Kantian perspective. In Bokulich, P. and Bokulich, A., editors, *Scientific Structuralism*, pages 1–24. Dordrecht, Heidelberg: Springer.

Maudlin, T. (2011). *Quantum Non-Locality and Relativity. Metaphysical Intimations of Modern Physics*. Malden, Oxford: Wiley-Blackwell, third edition.

Maudlin, T. (2014). Critical study: David Wallace, the emergent multiverse: Quantum theory according to the Everett interpretation. *Noûs*, 48(4):794–808.

Mermin, N. D. (1998). What is quantum mechanics trying to tell us? *American journal of physics*, 66(9):753–767.

Mittelstaedt, P. (2009). Cognition versus constitution of objects: From Kant to modern physics. *Foundations of Physics*, 39(7):847–859.

Muller, F. A. (2014). Elementary particles and metaphysics. In Galavotti, M. C., Dieks, D., Gonzalez, W. J., Hartmann, S., Uebel, and Weber, M., editors, *New Directions in the Philosophy of Science*, pages 417–431. Cham, Heidelberg: Springer.

Nagel, J. (2000). The empiricist conception of experience. *Philosophy*, 75(293):345–376.

Näger, P. M. (2016). The causal problem of entanglement. *Synthese*, 193(4):1127–1155.

Nakahara, M. (2003). *Geometry, topology and physics*. Bristol, Philadelphia: IOP Publishing Ltd, second edition.

Norsen, T. (2016). Quantum solipsism and non-locality. In Bell, M. and Gao, S., editors, *Quantum Non-Locality and Reality*, pages 204–237. Cambridge University Press.

Omnès, R. (1994). *The Interpretation of Quantum Mechanics*. Princeton University Press.

Petersen, A. (1963). The philosophy of Niels Bohr. *Bulletin of the Atomic Scientists*, 19(7):8–14.

Portmann, S. and Wüthrich, A. (2007). Minimal assumption derivation of a weak Clauser-Horne inequality. *Studies in History and Philosophy of Modern Physics*, 38:844–862.

Redhead, M. (2003). The interpretation of gauge symmetry. In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pages 124–139. Cambridge, New York: Cambridge University Press.

Redhead, M. L. G. (1975). Symmetry in intertheory relations. *Synthese*, 32(1/2):77–112.

Reichenbach, H. (1920). *Relativitätstheorie und Erkenntnis Apriori*. Berlin: Springer.

Rosenberg, J. F. (2005). *Accessing Kant*. Oxford, New York: Oxford University Press.

Ruetsche, L. (2011). *Interpreting Quantum Theories*. Oxford, New York: Oxford University Press.

Schlosshauer, M. (2007). *Decoherence and the Quantum to Classical Transition*. Berlin, Heidelberg: Springer, second edition.

Schrader, G. (1951). The transcendental ideality and empirical reality of Kant's space and time. *The Review of Metaphysics*, 4(4):507–536.

Schrödinger, E. (1935). Discussion of probability relations between separated systems. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):555–563.

Schwartz, M. (2014). *Quantum Field Theory and the Standard Model*. Cambridge, New York: Cambridge University Press.

Stanford, K. P. (2006). *Exceeding Our Grasp. Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.

Streater, R. F. (1988). Why should anyone want to axiomatize quantum field theory? In Brown, H. R. and Harré, R., editors, *Philosophical Foundations of Quantum Field Theory*, pages 137–148. Oxford: Clarendon Press.

van Fraassen, B. C. (1991). *Quantum Mechanics: An Empiricist View*. Oxford, New York: Oxford University Press.

van Fraassen, B. C. (2006). Structure: Its shadow and substance. *The British Journal for the Philosophy of Science*, 57(2):275–307.

van Fraassen, B. C. (2010). *Scientific representation: Paradoxes of perspective*. Oxford University Press.

von Borzeszkowski, H. and Mensky, M. B. (2000). EPR effect in gravitational field: nature of non-locality. *Physics Letters A*, 269(4):197–203.

Wallace, D. (2012a). Decoherence and its role in the modern measurement problem. *Philosophical Transactions of the Royal Society A*, 370(1975):4576–4593.

Wallace, D. (2012b). *The Emergent Multiverse. Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.

Wallace, D. (2020a). On the plurality of quantum theories: Quantum theory as a framework, and its implications for the quantum measurement problem. In French, S. and Saatsi, J., editors, *Scientific Realism and the Quantum*, pages 78–102. Oxford University Press.

Wallace, D. (2020b). The quantum revolution in philosophy. *Analysis*, 80(2):381–388.

Weatherall, J. O. (2016). Fiber bundles, Yang–Mills theory, and general relativity. *Synthese*, 193(8):2389–2425.

Weatherall, J. O. (2018). Regarding the 'hole argument'. *The British Journal for the Philosophy of Science*, 69(2):329–350.

Wiesing, L. (2014). *The Philosophy of Perception: Phenomenology and Image Theory*. Bloomsbury Publishing.

Wigner, E. (1939). On unitary representations of the inhomogeneous Lorentz group. *Annals of mathematics*, 40(1):149–204.

Wood, C. J. and Spekkens, R. W. (2015). The lesson of causal discovery algorithms for quantum correlations: Causal explanations of Bell-inequality violations require fine-tuning. *New Journal of Physics*, 17(3):033002.

Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica*, 43(1-2):99–124.

Worrall, J. (2020). Structural realism: The only defensible realist game in town? In Gonzales, W. J., editor, *New Approaches to Scientific Realism*, pages 169–205. de Gruyter.

Wüthrich, A. (2014). Local acausality. *Foundations of Physics*, 44:594–609.

Zurek, W. H. (1982). Environment-induced superselection rules. *Physical Review D*, 26(8):1862.

Zurek, W. H. (2005). Probabilities from entanglement, Born's rule $p_k = |\psi_k|^2$ from envariance. *Physical Review A*, 71(5):052105.