The Coalescence Approach to Inequivalent Representation: Pre-QM_∞ Parallels

Caspar Jacobs

Accepted for publication by the *British Journal for the Philosophy of Science* on 21-03-2021. DOI: 10.1086/715108.

Abstract

Ruetsche [2011] argues that the occurrence of unitarily inequivalent representations in quantum theories with infinitely many degrees of freedom poses a novel interpretational problem. According to Ruetsche, such theories compel us to reject the so-called 'ideal of pristine interpretation'; she puts forward the 'Coalescence Approach' as an alternative. In this paper I offer a novel defence of the Coalescence Approach. The defence rests on the claim that the ideal of pristine interpretation already fails before one considers the peculiarities of QM_∞: there are pre-QM_∞ parallels to coalescence. Despite this departure from pristinism, the 'modest' view that emerges poses no threat to scientific realism.

Contents

1 Introduction

In *Interpreting Quantum Theories* [2011], Laura Ruetsche argues that the occurrence of unitarily inequivalent representations in quantum theories with infinitely many degrees of freedom poses a novel interpretational problem. According to Ruetsche such theories, jointly labelled 'QM$_\infty$', compel us to reject the so-called 'ideal of pristine interpretation', which draws a sharp distinction between laws and initial conditions. Ruetsche puts forward her 'Coalescence Approach' as an alternative. Although this proposal has been recognized as standing at the centre of Ruetsche's book,[1] it has not received much attention in the literature.[2] Instead, discussions have focused on the relative merits of various 'pristine' interpretations of QM$_\infty$.[3] In this paper I offer a novel defence of the Coalescence Approach. The defence rests on the claim that the ideal of pristine interpretation already fails before one considers the peculiarities of QM$_\infty$: there are pre-QM$_\infty$ parallels to coalescence, from classical and statistical mechanics to 'ordinary' QM. I thus propose to extend Ruetsche's criticism of the pristine ideal (which I detail below) by drawing attention to the problems it faces in these different contexts.

I do so by distinguishing between a 'modest' and a 'radical' version of the Coalescence Approach. I argue that the modest version of the Coalescence Approach suffices both for the pre-QM$_\infty$ examples I discuss and for one of Ruetsche's main case studies, the phase structure of the ferromagnet. The significance of this distinction lies in its consequences for scientific realism. Ruetsche claims that the Coalescence Approach invalidates the No Miracles Argument for scientific realism. I point out that this is only the case if

[1] See reviews by Friederich [2013]; Wallace [2014].

[2] Lupher [2008] and Baker [2016] are the only notable exceptions that I know of.

[3] For example, Kronz and Lupher [2005] defend so-called 'Universalism', of which Baker [2009, 2016] favours a 'tempered' version. Feintzeig [2016, 2018a,b] offers various arguments in support of 'Algebraic Imperialism'. I am not aware of any advocates of 'Hilbert Space Conservatism'.

one adopts the radical Coalescence Approach. Its more modest version is compatible with realism. Since the latter suffices for all of the cases discussed in this paper, the demise of the pristine ideal poses no threat to realism.

The plan is as follows. In §2, I summarize the ideal of pristine interpretation. In §3, I explicate Ruetsche's HARMONY principle, which sets the standards for successful theory-interpretation. §4 recaps the threat that QM_∞ poses to the pristine ideal. In §5, I detail Ruetsche's alternative, the Coalescence Approach. In §6, I draw the distinction between the modest and radical version of the Coalescence Approach. Then, in §7, I discuss pre-QM_∞ cases of coalescence, drawn from classical mechanics, classical statistical mechanics and ordinary quantum mechanics. §8 connects these issues to the status of scientific realism, and §9 concludes.

2 The Ideal of Pristine Interpretation

Following Bas van Fraassen and Ruetsche, I take it that an interpretation answers the questions: '[u]nder what conditions is this theory true? What does it say the world is like?' (Van Fraassen [1991, 242]; see also Ruetsche [2011, 7]). Ruetsche explicates the process as follows.[4] In the 'syntactic stage', an interpreter specifies the mathematical states and observables. In the 'semantic stage', she assigns specific physical quantities to the observables, and hence constructs a map from the set of states into the space of possible worlds. The image of this map determines the physical possibilities.[5] Ruetsche's notion of a 'kinematic pair' is helpful here: a kinematic pair has the form $(\mathcal{S}, \mathcal{O})$, where $\mathcal{S}$ is a set of states and $\mathcal{O}$ a set of observables. Ruetsche's (controversial) core claim is that no single kinematic pair can capture the content of QM_∞.

An interpretation cleaves the space of possible worlds in two: some worlds are physically possible, others are physically impossible. An interpretation is 'pristine' when this cleaving satisfies the following conditions:[6]

1. UNIMODALITY. There is a single concept of physical possibility that applies equally to all states in $\mathcal{S}$;

[4]This account is slightly simplified; see Ruetsche [2011, Ch. 1] for details.

[5]There are two sorts of possibilities that states can represent. On one option, they represent possible instantaneous states of the world. On the other, they represent possible world-histories. There is no reason to categorically prefer one option over the other. In this paper, I will consider states which represent instantaneous possibilities and mostly set aside dynamical considerations.

[6]Although Ruetsche does not explicitly state these conditions, they are based on her description of pristinism in Ruetsche [2011, 4].

2. LOFTINESS. $(\mathcal{S}, \mathcal{O})$ is independent of (a) contingencies and (b) contextualities.

An interpretation is UNIMODAL when there is an unambiguous distinction between the physically possible and the physically impossible: there is no room for 'degrees' of possibility. LOFTINESS imposes the requirement that interpretations are independent of what Ruetsche calls 'geographical factors', such as initial conditions or particular applications. I will discuss the difference between contingencies and contextualities in more detail below.

What drives the ideal of pristine interpretation is a distinction between laws and initial conditions:

> The class of what applies in all settings where T applies includes T's laws [...]; the class of what changes from setting to setting includes initial/boundary conditions, as well as practical considerations parochial to the settings which give rise to them. To adhere to the ideal of pristine interpretation is to invoke only considerations from the former class when circumscribing the collection of worlds that are possible according to T. [Ruetsche 2011, 4]

The laws here include non-dynamical relationships, such as commutation relations. This distinction between laws and initial conditions depends on a clear concept of what remains constant across physically possible worlds. When the ideal of pristine interpretation is violated, this distinction becomes blurred.

3 The Harmony Principle

Ruetsche argues that pristine interpretations of QM_∞ fail. But in order to know whether an interpretation fails, we need to know what its success conditions are. It is therefore worthwhile to discuss these conditions in some detail.

Ruetsche bases her success conditions on a principle she calls HARMONY:

> HARMONY. $(\mathcal{S}, \mathcal{O})$ is an admissible kinematic pair for a theory only if (i) to each lawlike relationship posited by the theory there correspond elements of $\mathcal{O}$ standing in that relationship, and (ii) $\mathcal{S}$ includes all and only states instantiating the relationships in (i). [Ruetsche 2011, 143]

Ruetsche states Harmony as one principle, but it is useful to break it down into its constituent components. This yields the following set of conditions:[7]

1. O-completeness: If some physical observable O is physically relevant, then there is an element of $\mathcal{O}$ that represents O.

2. S-completeness: If a relation L between elements of $\mathcal{O}$ is a law, then all states in $\mathcal{S}$ instantiate L;

3. S-soundness: If all elements of $\mathcal{S}$ instantiate L, then L is a law.

The point of O-completeness is clear: it guarantees that there are enough observables to express the physically relevant facts, which includes the theory's laws. S-completeness and S-soundness express a relation between lawhood and necessity. The idea of S-completeness is that laws are physically necessary and so hold true in all physically possible states. Conversely, if there is a state in which some law-like relationship does not hold, then the latter does not count as a law. S-soundness, on the other hand, claims that the laws are laws in virtue of the fact that they hold true across all states, and hence if all physically possible worlds instantiate some fact, then that fact is *ipso facto* a law. The intuition here is that contingent generalizations such as 'there are no golden mountains' aren't laws exactly because there is some physically possible world in which there are golden mountains.

It is tempting to also posit the converse of O-completeness:

4. O-soundness: If an element of $\mathcal{O}$ represents a physical observable O, then O is physically relevant.

The motivation for O-soundness is to arrive at a minimal set of observables, which contains only those observables that are relevant for the characterization of the physical system in question. However, O-soundness is not entailed by Harmony, and there are a couple of questions one may have. Firstly, it isn't entirely clear when an observable is physically relevant—and without a clear account of such relevance, it seems presumptuous to use O-soundness to remove observables from our theories. Secondly, it is not obvious that it is a theoretical vice to have irrelevant observables, since those observables

[7] My terminology is meant to evoke principles of logical soundness and completeness: the mathematical relations between observables on mathematical states are 'syntax', whereas the relations between physical quantities in possible worlds are 'semantics'. S-soundness says that we can move from facts about the former to facts about the latter, and S-completeness says that we can move in the reverse direction.

could simply exist as a result of the overall mathematical structure of our theory. For example, the existence of a momentum observable $\hat{p}$ and a position observable $\hat{x}$ entails that there exists an observable $\hat{p} + \hat{x}$. But this latter observable seems physically uninteresting; it is only there because the addition of operators on a Hilbert space is well-defined. For these reasons, I will only use O-soundness as giving us a provisional, but not decisive, reason against having available certain observables.

4 Phase Structure: Threat to Pristinism

In this section, I will outline Ruetsche's main argument against the ideal of pristine interpretation, the 'Coalesced Structures Argument' [Ruetsche 2006, 2011]. Ruetsche's claim is that no pristine interpretation of quantum statistical mechanics can account for the phenomenon of phase structure. This motivates her preferred alternative, the Coalescence Approach. Let me emphasize that all of the arguments discussed below (except, perhaps, the one against Hilbert Space Conservatism) are to some extent controversial. It is not my aim here to defend all of Ruetsche's claims; rather, the purpose of this section is merely to motivate the consideration of the Coalescence Approach as an alternative to the pristine interpretations that Ruetsche criticizes.

The Coalesced Structures Argument focuses on phase structure. Phase structure is exemplified by ferromagnets, such as iron, which exhibit spontaneous magnetization below a certain critical temperature. We can model the ferromagnet as a spin chain by defining operators $\sigma^n_{i,j,k}$ that represent the canonical anti-commutation relations (CARs). The Jordan-Wigner theorem guarantees that for a finite system, this representation is unique up to unitary equivalence, so we can consider different representations of the CARs as mere notational variants (cf. Clifton and Halvorson [2001]). Now, a statistical system with finite DOFs has a unique equilibrium state.[8] However, the equilibrium state of the Heisenberg model of the ferromagnet is degenerate, due to the rotational symmetry of its Hamiltonian. Therefore, we must represent the ferromagnet as an infinite spin chain. In this context, the Jordan-Wigner theorem fails. For example, one representation of the CARs starts with a basis state in which the spin at all sites is +1 along the z-direction, and then adds further basis states which differ from the first at finitely many sites. Call this representation $(\mathcal{H}^+, S^+)$. But this represen-

[8] Where this is understood as a state that satisfies the so-called KMS condition; see Ruetsche [2003, §3] for details.

tation is not unique: we could equally well have started with a basis state in which all spin values are -1 along the z-direction. We then get an alternative representation of the CARs, $(\mathcal{H}^-, S^-)$. Crucially, $(\mathcal{H}^+, S^+)$ and $(\mathcal{H}^-, S^-)$ are unitarily inequivalent representations (UIRs). This presence of UIRs spells trouble for pristine positions.

4.1 Hilbert Space Conservatism

The first pristine position we discuss is 'Hilbert Space Conservatism', which claims that all quantum theories are Hilbert space theories: states are identified with positive, normalized trace-class operators $\mathcal{T}_1^+(\mathcal{H})$ on an irreducible Hilbert space representation $\mathcal{H}$ of the CARs, and observables are identified with bounded, self-adjoint operators $\mathcal{B}_{SA}(\mathcal{H})$. Faced with UIRs, the Hilbert Space Conservative has to choose one sole representation of the CARs as the bearer of physical content. The problem is that this violates S-soundness, which says that only the laws hold true across all possible states. Since all states on an irreducible representation of the CARs are polarized in the same direction, Hilbert Space Conservatism implausibly implies that it is a law-like fact in which direction a ferromagnet is polarized. Ruetsche puts it well: z$(\mathcal{H}^-, S^-)$ and $(\mathcal{H}^+, S^+)$ constitute distinct and rival theories of the infinite spin chain: the set of worlds possible according to the first theory includes none where the polarization is directed in the positive z direction; the set of worlds possible according to the second theory includes *only* worlds where the polarization is so-directed. Supposing both sorts of polarization are possible, neither $(\mathcal{H}^+, S^+)$ nor $(\mathcal{H}^-, S^-)$ can give the theory of the infinite spin chain' [Ruetsche 2006, 480].

4.2 Algebraic Imperialism

According to a second approach, 'Algebraic Imperialism', observables are self-adjoint elements A of a C*-algebra $\mathfrak{A}$, and states are linear maps $\omega : \mathfrak{A} \mapsto \mathbb{C}$ such that $\omega(A)$ is the expectation value of A in state ω. The algebra $\mathfrak{A}$ delineates a structure that is shared by all Hilbert space representations, and so $\mathcal{S}_{\mathfrak{A}}$ (the set of states over $\mathfrak{A}$) does contain states polarized in different directions. However, Ruetsche argues that Algebraic Imperialism violates O-completeness, the requirement that any physical quantity that figures in the laws is amongst the theory's observables. The reason is that there are nets of observables that converge in the weak operator topology of $\mathcal{H}$, but not in the norm topology natural to $\mathfrak{A}$; such observables are 'parochial' to particular Hilbert spaces. Crucially, the polarization observable $\hat{\boldsymbol{m}}$ is one of these parochial observables. But the phenomenon of phase structure is characterized through macroscopic observables such as $\hat{\boldsymbol{m}}$: in the paramagnetic phase $\hat{\boldsymbol{m}} = 0$, while in the ferromagnetic phase $\hat{\boldsymbol{m}} \neq 0$. Polarization observables are also essential to characterize the lawlike phenomenon of universality: the description of phase transitions in terms of critical exponents. Algebraic Imperialism seems unable to account for these aspects of phase structure. However, there are some possible responses. For instance, Feintzeig [2018a] has recently argued that Algebraic Imperialism does have access to parochial observables when extended to $\mathfrak{A}$'s bidual $\mathfrak{A}^{**}$. So whether Ruetsche's argument indeed succeeds remains the subject of debate.

4.3 Universalism

'Universalism' proposes a more inclusive representation: the universal enveloping Von Neumann algebra, $\pi_U(\mathfrak{A})''$, which is the weak closure of the direct sum π_U of the Hilbert space representations of all $\mathfrak{A}$-states (via the so-called GNS construction). Since states over $\mathfrak{A}$ are represented as normal states on $\pi_U(\mathfrak{A})''$, the Universalist has access to the same states as the Algebraic Imperialist. Furthermore, Feintzeig [2018a] has shown that there are observables on $\pi_U(\mathfrak{A})''$ that correspond to the parochial observables. However, many of these observables are physically irrelevant to characterize the actual state of a system, and so Universalism seems to fail O-soundness. Baker [2016] expresses (but does not endorse!) this point:

> Consider any two inequivalent, irreducible subrepresentations of the large representation. The parochial observables native to one subrepresentation will provide no physical information about the states of the other subrepresentation's folium. So these parochial observables play no role in describing the vast majority of the theory's states. And since every irreducible subrepresentation has its own family of parochial observables, that is a lot of surplus structure!

Recall that O-soundness motivates us to find a minimal set of observables that characterize some physical system. The problem for Universalism is that parochial observables on different subrepresentations seem irrelevant to the system's actual state, whereas a restriction to the irreducible representation of that state gives us just those parochial observables that are physically relevant for that state. Now, it may seem as if the universal representation also has a physically relevant observable native to it, namely the universal polarization observable—which we can think of as the sum of all parochial polarization observables. But in fact, it is only the projection of this observable onto the subrepresentation of the system's actual state that is relevant to the system's actual state , and this just is that subrepresentation's parochial polarization observable. On the other hand, the universal polarization observable's projection onto other subrepresentations does not tell us anything about the system of interest, so there is a sense in which Universalism doesn't give us a minimal set of actually relevant observables.

As I mentioned in §3, it remains controversial whether O-soundness indeed captures a genuine notion of physical redundancy. Moreover, even if Universalism fails O-soundness it is unclear how much of a problem this really is: surely too many observables are still better than too few? So I am aware that the above arguments has serious limits. In any case, Ruetsche has another argument, which is perhaps more persuasive: the 'W*-argument'. The upshot of the W*-argument is that Universalism fails S-completeness, which says that all states instantiate the laws. For there are certain Hamiltonians which fail to converge in the norm topology of $\mathfrak{A}$, but which do converge in the weak operator topology of some states' GNS representation. However, this convergence only holds for some states over $\mathfrak{A}$. If the Universalist considers all states over $\mathfrak{A}$ as physically possible, she includes states that don't admit of a dynamical evolution. The dynamics therefore wrongly fail to qualify as lawlike. In order to be more selective one has to consider 'geographical' factors that pertain to the particular dynamics of the system, contrary to LOFTINESS. But the W*-argument is also open to critique. For example, Feintzeig [2018b] discusses methods for constructing alternatives C*-algebras that do satisfy S-completeness. As with Algebraic Imperialism, then, these arguments at best motivate a consideration of Ruetsche's alternative approach, without offering decisive reasons in its favour.

5 The Coalescence Approach

Ruetsche's verdict is clear: the interpretative strategies discussed above falter because of their pristinism. Each decides on a kinematic pair '*a priori*, before the messy business of applying the theory in question to individual problems begins' [Ruetsche 2011, 146]. We can diagnose the issue as follows: S-soundness and O-completeness pull in the direction of admitting more states and observables, whereas S-completeness and O-soundness push us towards fewer of both. However, which states and observables are relevant depends on 'geographical factors'. Pristine interpretations either ignore those factors entirely or focus on a particular set of factors *a priori*. We need an approach that balances these demands.

Ruetsche puts forward an unpristine alternative:

> 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. Instead of one possibility space pristinely associated with a theory from the outset, many different possibility spaces, keyed to and configured by the many settings in which the theory operates, pertain to it. Following Kadison, I call this the *coalescence approach* to interpreting physical theories. [Ruetsche 2011, 147]

As it stands, this Coalescence Approach (CA) needs further spelling out. The remainder of this section offers one such account.[9]

Recall that pristine interpretations put forward a single kinematic pair that carries a theory's content. By contrast, the CA has an (at least) two-tiered account of possibilities. The first tier consists of a sole, primitive kinematic pair which represents a broad space of physical possibilities. The second tier consists of a 'kaleidoscope' of kinematic pairs, each indexed to one or more geographical factors. Because of its two-tieredness, the CA violates UNIMODALITY, which holds that there is a unique notion of physical possibility that applies equally to all states. Moreover, the CA also violates LOFTINESS, which demanded that such factors play no role in theory-interpretation, because the second tier possibilities are indexed to geographical factors. So the CA radically departs from the pristine ideal.

The key idea behind Ruetsche's Coalescence Approach, as I understand it, is that there is a 'division of labour' between the two tiers such that each tier fulfils distinct success conditions. The broad kinematic pair at Tier 1 satisfies O-completeness and S-soundness: it contains all observables one could need, and enough states to avoid contingent generalizations. The elements of Tier 2, on the other hand, satisfy O-soundness and S-completeness: they contain a minimal set of relevant observables, and few enough states to sustain the laws. The Coalescence Approach aims to avoid the problems that beset the ideal of pristine interpretation through this division of

[9]For another attempt at spelling out Ruetsche's position, see Lupher [2008, §5.6]. I borrow some elements from his set-up, including the Tier 1/Tier 2 language.

labour. Rather than fulfilling all success conditions at once, the Coalescence Approach delegates different duties to different kinematic pairs.

In more detail, Ruetsche recommends that we allow Tier 1 to range over the space of all algebraic states, $\mathcal{S}_{\mathfrak{A}}$, which represent 'the worlds possible according to the theory in the broadest sense' [Ruetsche 2011, 290]. In an earlier paper Ruetsche [2003] suggests that we can also use the universal representation $\pi_U(\mathfrak{A})''$ at Tier 1. The normal states on $\pi_U(\mathfrak{A})''$ are identical to the states over $\mathfrak{A}$, so on that score the difference is irrelevant. But the latter has access to parochial observables which the former lacks, and this is significant. *Pace* Ruetsche, I believe that it is essential to have these observables available at Tier 1. Recall that the reason for preserving the full set of states $\mathcal{S}_{\mathfrak{A}}$ at Tier 1 is to ensure the truth of counterfactuals such as: 'if the external magnetic field had had a different value, the ferromagnet would have been polarized in a different direction'. As it stands, $\mathcal{S}_{\mathfrak{A}}$ lacks the resources to express those counterfactuals, because they involve parochial observables. On the other hand, Feintzeig [2018a] proves that there are operators on the universal representation that correspond to the parochial observables [Feintzeig 2018a, Prop. 5]. Their expectation value is as expected on their 'own' subrepresentation, and zero otherwise. Furthermore, Feintzeig argues that a reasonable extension of Algebraic Imperialism to the bidual $\mathfrak{A}^{**}$ allows the latter access to the same observables (Prop. 2). If we accept this extension, Imperialism and Universalism are in fact equivalent up to *-isomorphisms (Prop. 3). Therefore, I recommend that the appropriate Tier 1-pair is either the universal enveloping Von Neumann algebra or the C*-algebra as extended in the way Feintzeig suggests.

Tier 2, on the other hand, 'does take contingencies into account'. From $\mathcal{S}_{\mathfrak{A}}$ or $\pi_U(\mathfrak{A})''$ one selects a more narrow space of states 'relevant to the application at hand'. As Ruetsche writes, 'other algebraic states aren't impossible; they're simply possibilities more remote from the present application of the theory [...]' [Ruetsche 2003, 1340]. While focusing on a particular space of states, one also 'coalesces' a set of observables with them. For example, we can construct the GNS representation of the ferromagnet's actual state ω, close in the weak operator topology and so obtain the relevant parochial observables. Again, 'observables confined to von Neumann algebras parochial to less relevant representations aren't once and for all unphysical. They're just inadequate to the sorts of discriminations demanded by the application at hand' [Ruetsche 2011, 290]. In this sense the GNS representation coalesces a 'minimal' set of observables most relevant to the particular state of the system. There are two reasons for this restriction. The first is to describe the system of interest in the simplest terms possible: the CA sees simplicity as a theoretical virtue (cf. Ruetsche's [2011, §15.4] invocation of Humean accounts of laws). The second is to allow for certain explanations of phenomena in terms of natural laws: sometimes a relation only appears lawlike from the space of Tier 2 states, so an interpreter who appeals to that relation's physical necessity requires this more restricted possibility space. In this way, the Coalescence Approach supports a theory's explanatory aspirations.

In sum, the lenient representation at Tier 1 allows us to claim that distinct polarizations are physically possible, whereas the restricted representation at Tier 2—sensitive to *a posteriori* factors—hones in on the states and observables most relevant to the actual state of the ferromagnet. On my reading of the Coalescence Approach, then, it is this division of labour that allows us to avoid the choice of a single kinematic pair that either has too few states and/or observables, or contains non-dynamical states and/or irrelevant observables.

6 Modest and Radical

The 'geographical' factors that determine a Tier 2-representation come in two kinds. The first are contingencies, or initial conditions. Roughly, initial conditions vary across worlds but are constant within them, whereas context also varies within a possible world. The latter includes the goals, applications and aspirations of scientists. We can thus draw a distinction between a 'modest' and a 'radical' CA:

> Modest CA: The interpretation of a theory depends on contingent facts about the target system's state, but not on the context in which the theory is used.
>
> Radical CA: The interpretation of a theory depends on contingent facts about target system's state and on the context in which the theory is used.

On the modest CA, the set of states $\mathcal{S}$ depends on initial conditions: possibility is indexed to possible worlds. For example, whether a certain ferromagnet is in its paramagnetic state at time t differs across worlds, but is fixed within each world. Therefore, the appropriate Tier 2-representation of a ferromagnet is fully determined by its actual state. On the radical CA, on the other hand, the state of a system does not fully determine its interpretation. Instead, the same system admits of different interpretations—different kinematic pairs—that depend on factors external to the system itself. Of course, there is a sense in which a system's state itself may depend on external factors. For example, a scientist may purposefully prepare a ferromagnet in its + state with some particular application in mind. However, the claim that interpretations depend on context is not a causal but a supervenience claim. Whatever it was that caused the ferromagnet's actual state, the modest CA demands that its kinematic pair fully supervenes on that state, whereas the radical CA allows that other facts play a role in the determination of what is physically possible to that system.

In Ruetsche's terms, the modest CA naturalizes physical possibility, whereas the radical CA also pragmatizes physical possibility:

> When physical possibility is naturalized, sets of physically possible worlds are indexed with respect to the 'anchor' world whose facts shape that set. When physical possibility is pragmatized, sets of possible worlds are indexed (or indexed as *well*) to circumstances of application within the 'anchor' world. When physical possibility is pragmatized, there's a single way the world is, but (as gleaned by physics) there isn't a single set of ways it might be. [Ruetsche 2011, 353]

As the above quote illustrates, Ruetsche is sensitive to these differences, but she does not explicitly distinguish between the above two versions of the CA. It is clear that Ruetsche would prefer its radical version: 'In slogan form, my contention is that, when it comes to QM_∞ at least, there is a pragmatic dimension to theory articulation. What set of possible worlds we associate with a theory of QM_∞ can depend on what we'd like to do with that theory: what explanations, involving which magnitudes and guided by what laws, we aspire to; what phenomenological models we need to construct; what projects of theory development we'd like to sponsor.' [Ruetsche 2011, 352].[10] Nevertheless, it is not always clear whether Ruetsche's evidence for the Coalescence Approach supports just the modest or also the radical version. For example, the coalescence account of the ferromagnet only requires that interpretations vary with initial conditions. As I explained in the previous section, the key idea is that a ferromagnet's Hilbert space representation at Tier 2 depends on its actual state, viz. its direction of polarization. If the ferromagnet is in the + state, for example, its kinematic pair at Tier 2 is $(\mathcal{H}^+, \mathcal{S}^+)$. There is no further need for contextual considerations. This is

[10] In personal communication, Ruetsche has confirmed that she identifies as a radical coalescer.

relevant because the modest and radical CA seem to have different implications for scientific realism. Specifically, I will argue in §8 below that the the former is consistent with realism, *pace* Ruetsche.

7 Pre-QM$_\infty$ Coalescence

I now come to the main claim of this paper: that there are analogues of the (modest) Coalescence Approach for pre-QM$_\infty$ theories. I will consider three examples in this section. The first derives from classical mechanics and concerns the choice of an N-dimensional phase space. The second example constructs the Past Hypothesis in statistical mechanics as a case of coalescence. The third and final example considers unitarily inequivalent representations in 'ordinary' quantum theories (that is, quantum theories with finitely many DOFs), illustrated by the Aharonov-Bohm effect. In each instance, my claim is that the way we interpret these theories runs counter to the ideal of pristine interpretation. The Coalescence Approach, meanwhile, aptly accounts for our interpretational practices. *Pace* Ruetsche, then, I claim that the tensions in the pristine ideal are already apparent before one considers QM$_\infty$. Moreover, we will see that in each case the modest CA suffices to account for the way we interpret these theories, an observation that will become relevant in the next section. Let me stress that the below examples do not rule out the radical CA, but rather suggest that appeals to context are not required in order to satisfy the adequacy conditions on interpretation encoded in Harmony via a 'division of labour'. As I will show in the next section, this means that these examples of pre-QM$_\infty$ coalescence do not immediately threaten realism.

7.1 Classical coalescence: phase spaces

In classical mechanics, we represent a system of n point-like masses in three-dimensional space on a $6n$-dimensional phase space, isomorphic to $\mathbb{R}^{6n}$: points of phase space correspond to possible states of the system. I will sketch some pristine interpretations of classical phase spaces below, but I warn from the outset that these positions may seem trivial and perhaps even silly. But that's exactly the point: pristine interpretations of classical mechanics are 'extremist' and ill-suited to physical practice. The classical analogue of the Coalescence Approach, on the other hand, offers a natural account of our use of classical phase spaces.

As our first extremist we can imagine a 'Phase Space Conservative' who, analogous to her quantum counterpart, privileges a single phase space amongst all others, for instance $\mathbb{R}^{6N}$ for some particular N. Since all states on $\mathbb{R}^{6N}$ represent classical worlds with N bodies, Phase Space Conservatism implies that those worlds are the only physically possible ones. But recall that S-soundness says that only the laws hold true across all states, and surely it is not a law that there are N bodies in the universe! Just as Hilbert Space Conservatism erroneously implied that any particular ferromagnet only could have been polarized in the direction in which it is actually polarized, Phase Space Conservatism cannot account for the fact that a classical universe could have contained either more or fewer bodies than it actually does.

Alternatively, we can envisage a 'Classical Universalist', who proposes that the disjoint union of phase spaces, $\bigsqcup_{n\in\mathbb{N}} \mathbb{R}^{6n}$, is the correct arena for classical mechanics (perhaps we could think of this space as a 'classical Fock space').[11] This space is an immensely complex mathematical structure: one can use it to represent classical states for any $n \in \mathbb{N}$. The problem with Classical Universalism is that it is extremely unparsimonious. Suppose, for example, that the actual universe contains ten bodies, so that its state lies on a subspace of the universal phase space isomorphic to $\mathbb{R}^{60}$. But the universal phase space also contains observables that are inapplicable to states on $\mathbb{R}^{60}$, such as the position of the eleventh particle or the momentum of the thirteenth. These observables are defined over $\mathbb{R}^{66}$, or $\mathbb{R}^{78}$. Therefore, Classical Universalism violates O-soundness, just as the (quantum) Universalism discussed in §4.3 does. But while the excess observables of the universal enveloping Von Neumann algebra were, perhaps, justifiable as by-products of otherwise physically meaningful structure, this is not the case here: there is no sense in which $\bigsqcup_{n\in\mathbb{N}} \mathbb{R}^{6n}$ is the 'natural' arena of classical mechanics as applied to the actual universe. Classical Universalism's redundancy of observables motivates an alternative approach.[12]

In practice, physicists simply use the phase space which is most applicable to the system at hand, but pristine interpretations cannot easily account for this fact.[13] The Coalescence Approach offers a better model of these

[11] I will not discuss classical analogues of Algebraic Imperialism. Feintzeig [2016] constructs an algebraic version of classical field theory, but this construction is irrelevant for our purposes as the focus of this subsection is on particle theories with varying N.

[12] Baker [2016] objects that this argument simply 'amounts to the observation that, on a universalist-style view, physically contingent facts have a significant role in determining which physical quantities [...] provide physically significant information'. But this misses the point, which is exactly that Universalism, unlike the CA, is not sensitive to the way in which contingent facts determine which observables are most significant.

[13] An anonymous reviewer has suggested that the pristine ideal can model physical practice in terms of idealizations. But it is not clear to me whether the restriction to

practises. The idea is that at Tier 1, the classical coalescer has phase spaces for all n at her disposal. This accounts for the fact that different particle numbers are physically possible, in accordance with S-soundness. The 'excess' observables are necessary at this level to express such counterfactuals. On the other hand, Tier 2 selects the particular phase space that matches the number of DOFs of the system under study. As a result, the kinematic pair at Tier 2 only contains observables that directly bear on said system. This representation thus satisfies O-soundness. This is the same division of labour that I described in §5. Therefore, the CA is a more natural account of our use of classical phase spaces, as it avoids both a too restrictive notion of physical possibility and a too liberal account of observables. Instead, it affirms the fact that physicists are not limited to the use of a single phase space, but choose opportunistically between them.

Moreover, this employment of the CA is modest in the sense that it only appeals to contingent facts about the number of DOFs of a particular system. Of course, the practising coalescer may still use different phase spaces to model different sub-systems of the universe. But this is not an example of context-dependence as outlined in §6, since in each case the appropriate choice of phase space is fully determined by the physical state of the system of interest. This is not to say that context could never play a role in theory-interpretation. I will discuss potential limits to my claim that the modest CA suffices in the next section.

Before I move on to the second example, let me consider an important disanalogy between the polarization observable $\hat{\mathbf{m}}$ and the particle number n. While n is a conserved quantity, $\hat{\mathbf{m}}$ is not—the latter, after all, varies in phase transitions. Therefore, one may object that while the same ferromagnet could have been polarized in different directions, it is not the case that the same system could have had a different number of DOFs. This is a difficult issue which seems to depend on intuitions: to me, it seems possible that, say, the solar system could have had one fewer planet while still being the same system. Similarly, our universe as a whole could have contained one less particle than it actually does, but still count as the same universe (or, at least, as our universe's other-worldly counterpart). Of course, there are limits: it is unlikely that the free particle could have been a tiger. And there is significant vagueness, too: Theseus' ship is (arguably) still Theseus' ship when one of its ores is removed, but perhaps no longer when the mast

---

a particular phase space amounts to an idealization: an idealization usually ignores or distorts some actual features of the system, whereas a particular choice of phase space dismisses some of the system's possible states.

and hull are replaced. Putting that aside, the fact that there are no dynamical transitions between different values for n does not establish that a particular value of n is essential to a system, in the *de re* sense that a system must have the same number of particles in all worlds in which it exists. Therefore, I believe that the cases are sufficiently analogous for the parallel to hold.[14]

7.2 Statistical mechanics: emergence as coalescence

A different example of the Coalescence Approach in pre-QM_∞ action is the Past Hypothesis in statistical mechanics. In brief, the fact that the laws of statistical mechanics are symmetric under time reversals seems to contradict the asymmetry of the Second Law of Thermodynamics, which states that entropy (probably) increases with time. In order to account for this asymmetry, the 'Past Hypothesis' (PH) is postulated: the entropy of the universe at the Big Bang was (sufficiently) low. This postulate ensures that the Second Law holds for future times [Albert 2000].

Since we are considering dynamics, it is useful to consider trajectories rather than points on phase space. Pristine interpretation then comes down to an *a priori* circumscription of allowed trajectories; the imposition of the PH in effect amounts to selecting only those trajectories with a low initial entropy as physically possible. But pristinism faces the following dilemma. On the one hand, if the Second Law is indeed a law, then S-completeness demands that all (or at least most) trajectories through phase space instantiate it: hence we impose the PH. On the other hand, the PH seems a contingent fact—the universe could have had a higher initial entropy.[15] So imposing the PH violates S-soundness, the principle that some fact is a

[14] There are also more realistic examples, which I lack the space to discuss in detail. One class of examples concerns the phenomenon of 'geometric phases'. The crux is that when some quantity μ, such as angular momentum, is conserved, it is often fruitful to treat μ as a parameter with a constant value c. One then constructs the 'reduced phase space' of a system, which contains just those points for which $\mu = c$. The dynamics on the reduced phase space P_μ non-trivially depend on the value of μ. As in the examples above, it is neither wise to assert the primacy of a unique reduced phase space, thereby ruling out counterfactual values of μ as physically impossible, nor to insist on the full structure of the universal phase space, which contains excess degrees of freedom with respect to the actual dynamics of the system. Instead, different systems call for different choices of P_μ, as the CA advocates. I thank an anonymous referee for suggesting this example.

[15] This is true, for example, on Lewis' [1973] version of the best systems account, which only counts generalizations as candidates for lawhood. On the other hand, Callender [2004] argues that the Past Hypothesis is a (non-dynamical) law. I won't further discuss this issue here.

law when all states instantiate it. Therefore, pristinism cannot fulfil both S-soundness and S-completeness simultaneously.

The Coalescence Approach avoids this dilemma with its familiar division of labour. Let Tier 1 consist of all possible trajectories, including those not instantiating the PH, and suppose that Tier 2 is indexed to particular trajectories on Tier 1. If the Second Law holds in the actual world, then Tier 2 includes only those trajectories which also instantiate the PH. From a Tier 1 perspective, then, the PH is a contingent fact, analogous to the polarization of the ferromagnet or the number of bodies in the universe. But since the PH holds true in our world, there are some trajectories which are 'more' physical than others. These low-entropy trajectories are coalesced at Tier 2, which thus represents the Second Law as a law. According to the CA, the Second Law is a law relative to our world: in a universe which starts out in equilibrium, for example, Tier 2 will also contain trajectories which fail to instantiate the Second Law. But this is correct, since the Second Law in a sense 'emerges' from the laws of statistical mechanics in combination with a particular low-entropy initial condition. Therefore, the Coalescence Approach offers a novel understanding of the reduction of thermodynamics to statistical mechanics. Moreover, which trajectories are most physical clearly depends solely on the initial condition of the actual world. It follows that here too the modest CA suffices to account for the Second Law.

As Baker [2016] points out, the Past Hypothesis is not the only instance of such emergence:

> [I]f one is convinced that the postulates [such as the PH] are laws, and that it is important in other explanatory contexts (for example, ones appealing to time-reversal symmetry) to suspend these postulates, then coalescence is everywhere in physics.

I agree with this assessment: the Coalescence Approach is essential to understanding pre-QM_∞ theories. In addition to thermodynamics, Baker mentions retarded electromagnetic waves; the time-asymmetric processes described in Wallace [2017], such as radioactive decay, allow a similar analysis. These considerations further confirm this section's point that the Coalescence Approach is not a radical departure from, but a familiar feature of scientific practice. The application of coalescence to other cases of emergence is a fruitful avenue for further research.

7.3 Quantum mechanics: bead on a wire

The final example is perhaps closest to the occurrence of UIRs in QM_∞. For an infinite number of DOFs is only one way in which the Stone-Von Neumann and Jordan-Wigner theorems can fail. These theorems also fail for finite systems whose configuration spaces are non-simply connected. I will discuss a rather simple example, namely the bead on a wire. But this example is a simplified version of the Aharonov-Bohm effect. Earman [2019, 2010] suggests that it 'would be interesting to compare this case [the Aharonov-Bohm effect] to the cases in QFT and the thermodynamic limit in models of phase transitions where there is also a choice among unitarily inequivalent representations'; this subsection carries out that comparison. The exposition follows Ruetsche [2013]; for connections to the Aharonov-Bohm effect, see Landsman [1990], Belot [1998] or Earman [2019].

Consider a bead constrained to move on a circular wire. In polar coordinates, its state is parametrized by the variables ϕ and ω, which represent the angular position and momentum respectively. However, ϕ is discontinuous at $\phi = 2\pi$, which hinders the process of quantization; a better choice of variables is $x = \cos\phi$, $y = \sin\phi$ and ω. Due to the constraint that $x^2 + y^2 = r^2$, the configuration space is just the circle S^1. Since $\omega \in \mathbb{R}$, the phase space of the bead on the wire then is $S^1 \times \mathbb{R}$. In order to see the connection with the Aharonov-Bohm effect, think of the bead as an idealized charged particle; its phase space is non-simply connected because the particle moves around an infinitely long impenetrable solenoid, which thus removes the z-axis from configuration space.

To quantize this theory, we require a set of commutation relations. These are what Ruetsche calls the 'Circular Canonical Commutation Relations' or CCCRs:

$$[x, y] = 0 \quad [\omega, x] = iy \quad [\omega, y] = -ix$$

In order to find a Hilbert space representation of the CCCRs, we need both a Hilbert space and operators that act on it satisfying these relations. The appropriate Hilbert space is $L^2(S^1)$, the space of square-integrable functions $\psi(\phi) : S^1 \to \mathbb{C}$. The (wave)function ψ assigns a complex number to each point of the bead's configuration space. As our observables we can choose:

$$\hat{x} = \cos\phi \quad \hat{y} = \sin\phi \quad \hat{\omega} = i\frac{d}{d\phi}$$

It is easy to show that these observables satisfy the CCCRs. However, this quantization is not unique. In fact, any choice $\hat{\omega}_\theta = i\frac{d}{d\phi} + \theta$, where

$\theta \in [0, 1]$, satisfies the CCCRs. In other words, any choice of θ yields a different representation of the CCCRs. Furthermore, these representations are unitarily inequivalent. The reason is that the spectra of $\hat{\omega}_\theta$ and $\hat{\omega}_{\theta'}$ for $\theta \neq \theta'$ are distinct. Specifically, the spectrum of $\hat{\omega}_\theta$ is $\{2\pi(n - \theta)\}$, where $n \in \mathbb{Z}$. We thence have a one-parameter family of UIRs.

In order to connect this to the Coalescence Approach, recall that $\hat{\omega}$ represents angular momentum. Therefore, different representations of the CCCRs deem different angular momentum values physically possible. This poses now-familiar problems for our pristine approaches. For Hilbert Space Conservatism, the physically possible values of angular momentum are limited to those within the spectrum of a single choice for $\hat{\omega}_\theta$. This limitation conflicts with our intuition that the bead could have had any angular momentum. Moreover, in the Aharonov-Bohm effect the value of θ depends on the magnetic flux through the solenoid: different values of θ result in empirically distinct interference patterns on the screen. Therefore, the Conservative's *a priori* restriction to a unique choice for θ rules out a whole class of experimentally detectable phenomena, in violation of S-soundness.

Universalism, on the other hand, claims that the direct sum of these representations is the correct setting for ordinary QM. This implies that angular momentum operators of all representations are physically significant, even though the actual value of $\hat{\omega}$ always lies within the spectrum of one of its representations. The excess of observables may violate O-soundness: just as polarization observables of different representations are irrelevant to the actual state of a ferromagnet, so different angular momentum observables are irrelevant to the actual momentum of the bead on the wire. Furthermore, there is a sense in which Universalism violates S-completeness, analogous to the W*-argument. Recall that one can represent the time-evolution of a quantum system either as a family of *-automorphisms on $\mathfrak{A}$, or via a Hamiltonian that acts on $\mathcal{H}_\theta$. The latter depends on the choice of representation for the angular momentum observable: different representations yield different Hamiltonians. This is the case for the particle in the Aharonov-Bohm effect [Landsman 1990; Earman 2019]. Now, consider a system with angular momentum $\hat{\omega}$. The dynamics for this state are given by $\hat{H}_\theta$, the particular Hamiltonian that acts on $\mathcal{H}_\theta$. But this Hamiltonian is parochial to the irreducible subrepresentation π_θ, and so there is a sense in which the universal representation contains dynamically irrelevant (or less relevant) states. Put differently, the Hamiltonian $\hat{H}_\theta$ corresponds to the projection of a unitary evolution operator $\hat{H}$ onto the subrepresentation $\mathcal{H}_\theta$. But it is the restriction of this Hamiltonian to $\mathcal{H}_\theta$ that we are ultimately interested in, so Universalism seems to give us more than we care about. In sum, the

Universalist cannot draw a distinction between those representations of the system's time-evolution that are dynamically distinguished and those that are not.[16] It is this distinction that requires us to attend to geographical factors.

The Coalescence Approach, meanwhile, advocates flexibility. The Tier 1-representation is the universal one, which allows for states of all possible angular momenta. The Tier 2-representation then hones in on a particular $\hat{\omega}_\theta$, and so sheds physically irrelevant content. This choice is not made once and for all, but depends on the actual state of the physical system. In the Aharonov-Bohm effect, for example, this choice is determined by the magnetic flux through the solenoid, which has observable consequences for the interference pattern. The CA avoids the issues discussed in the previous paragraph. Firstly, the CA coalesces a particular representation of $\hat{\omega}$ as physically most relevant, in accordance with O-soundness. Secondly, the CA also coalesces a particular representation of the dynamical evolution as most relevant to the actual system. In this sense, the CA offers us a better account of the relation between the dynamics at the algebraic and the representation level, since it allows for the fact that at Tier 2 there exists an empirically distinguished representation of the family of *-automorphisms at Tier 1. Finally, as with the previous examples, this is an instance of the modest CA: interpretation depends on contingent facts, such as the value of θ, but not on contextual factors. Again, my claim is not that this rules out the radical CA, but that at least in the first instance the radical CA is not required to account for the bead on the wire. As I show in the next section, this means that a coalescence account of ordinary QM is consistent with scientific realism. I will postpone a discussion of the potential limits of this claim to the same section.

8 Realism Restored

I have mentioned several times that what is at stake in this debate is the status of scientific realism. Specifically, Ruetsche [2011, §15.3] argues that the 'No Miracles Argument' (NMA) presupposes the ideal of pristine interpretations, so the Coalescence Approach suggests a move away from realism. However, I will argue this is at most a consequence of the radical CA.

[16] Neither can Algebraic Imperialism fully account for this distinction, since the Hamiltonian 'may depend explicitly on the parameter ω labelling the representation π_ω of the quantum algebra' [Landsman 1990, 15]. The distinctions between representations that Imperialism disavows are thus empirically highly significant.

The modest CA, on the other hand, is consistent with—indeed supports—scientific realism. Since the examples discussed above all seem to utilize the modest CA, I maintain that it is far from clear that the demise of the pristine ideal is a threat to realism.

The NMA argues that the success of our theories warrants belief in their (approximate) truth. Ruetsche rightly points out that the NMA concerns the success and truth of interpreted theories. After all, un-interpreted theories have no physical content: they do not 'say what the world is like'. Therefore, the truth of an (interpreted) theory is inferred from its success under the same interpretation. For example, if Everettian QM is successful in some way, then this supports the truth of QM under the Everett interpretation, but not the Bohmian one. For our purposes here, an interpretation consists of a pair of possibility spaces: one for each of the two tiers of the CA. So, what the realist is realist about are the notions of physical possibility that these spaces encode: whether a state of affairs is physically possible—in either the broad Tier 1 sense or the narrow Tier 2 sense—is supposed to be an objective matter of fact. I will argue that the modest version of the CA is consistent with the claim that these notions of physical possibility are objective.

On the Coalescence Approach, theories have no unique interpretation. Instead, interpretations may depend on geographical factors. Specifically, the possibility space at Tier 2 varies with context and/or initial conditions. Supposing that all these interpretations are roughly equally successful, the dilemma for realism is this: either each interpretation is individually successful enough to merit the status of approximate truth, or none are. In the first case, the NMA generates contradictory conclusions: it implies that the theory is true on each successful interpretation. But these interpretations disagree on which worlds are physically possible—that's what makes them distinct interpretations! Since the realist asserts that physical possibility is an objective notion, these contradictory commitments are worrisome. The second option is not much better: if geographically adulterated interpretations are not individually successful, the NMA finds no application at all. This means that there simply is no reason to believe that the theory's possibility spaces are objective. As Friederich [2013] notes: 'Given the importance and tremendous success of quantum theories, this challenge to scientific realism deserves serious consideration. Due to the clarity of Ruetsche's formulation of it the stage is well set for the realists to come up with their rejoinders.'

But whether Ruetsche's argument is valid depends on how we construe the CA. If we only accept the modest CA, the problem for scientific realism does not arise. The modest CA enables us to give a single (albeit adulterated) interpretation of QM_{∞}. This interpretation depends, as we have seen, on the initial conditions of the actual world. But all this means is that scientists in different worlds have different commitments, and this is no threat to realism. After all, realism claims that the success of our theories justifies the belief that our world is one of that theory's models. For example, realism about the coalescence account of classical phase spaces discussed in §7.1 implies two commitments: (i) a commitment to the physical possibility of different particle numbers in the 'broad' sense relevant to our modal intuitions, and (ii) a commitment to the physical necessity of the system's actual particle number in the 'narrow' sense relevant to the most simple characterization of its state. Realism is compatible with the modest CA in the sense that neither of these commitments varies within possible worlds, although the second commitment does vary across possible worlds.

On the other hand, the radical CA is indeed susceptible to Ruetsche's dilemma. According to this version of the CA there is no single interpretation of a theory to which all of its theoretical virtues accrue, even within our actual world. Instead, there are competing interpretations with distinct and conflicting metaphysical commitments. Which interpretation is (approximately) true depends on the context. In particular, it is context-dependent which states are physically possible in the narrow sense of Tier 2. Therefore, Tier 2-possibility is not an objective notion, but a mind-dependent construct. This presents a clear conflict with scientific realism. Now, it may seem as if this conclusion is an artifact of the way we have construed what counts as an interpretation of a theory. Suppose that an interpretation consisted not just of two possibility spaces, one of which varies with context, but of a possibly infinite collection of kinematic pairs, each of which is relevant in a different context. In that case, context does not determine what is physically possible: it merely directs our attention to a different notion of physical possibility. But the problem with this suggestion is that one cannot simply call a kinematic pair a possibility space; one also has to justify what particular notion of physical possibility this kinematic pair encodes. And, crucially, this justification cannot itself appeal to context, for then said notion of physical possibility is not sufficiently objective. For instance, I have couched the Tier 1/Tier 2 distinction in terms of the soundness and completeness principles of §3. Without a similar justification for their novel possibility concepts, the advocate of the radical CA remains committed to notions of physical possibility whose role in the theory depends on our subjective interests.

In conclusion, the modest CA is consistent with scientific realism. Since the examples of §7 only seem to require the modest CA, realism is safe from Ruetsche's threat. However, this conclusion clearly rests on the assertion that the modest CA is sufficient to account for pre-QM_∞ cases of coalescence. In closing, then, let me consider some limitations of this claim. Firstly, I want to stress that the modest CA does not exclude context from science altogether. It may very well be that our goals, applications and aspirations determine some particular possibilities as especially relevant to us. For example, suppose that I desire to study the effect of certain values of small variations of $\hat{\omega}$ around some fixed value $\hat{\omega}_0$. I may then consider the direct sum of representations for just those values of $\hat{\omega}$ close to $\hat{\omega}_0$. But this only contradicts the modest CA if it is additionally claimed that this sum corresponds to some novel notion of physical possibility. In that case, physical possibility depends on context. But there simply is no reason to believe that: the fact that we are interested in some particular representations does not elevate them to scientific concepts. So, the modest CA is not inconsistent with the commonplace claim that context plays some role in scientific representation.

But the above discussion does hint at a more worrisome scenarios for the coalescence realist, namely a scenario in which some notion of possibility that does carry metaphysical weight depends on context. For example, consider our representation of the solar system. If we are interested in the period of Jupiter around the Sun, it does not do much harm to idealize the heavenly bodies as point masses, so that we can represent the system on an (6×8)-dimensional phase space. But if we are interested in whether Jupiter and the Sun will collide, we must take their radii into account. This adds eight additional degrees of freedom to our theory, and hence we must use a (7×8)-dimensional phase space. The problem now is that both of these phase spaces lay claim to a notion of physical possibility that licenses certain scientific explanations, for example of Kepler's laws. But they disagree on which states are physically possible: in particular, the first phase space literally says that it is impossible for the planets to have a non-zero radius. So, what is physically possible (in the narrow sense) seems to depend on what we are trying to explain. The bad news is that I have no easy answer to this challenge, and so the practice of idealization—which seems essential to science—poses a serious threat to realist interpretations of the modest CA. But the good news is that idealization in any case raises difficult question about the relation between possibility and explanation. For example, Bokulich [2011] argues that idealized models only represent certain 'real patterns of structural dependencies in the world'. If that's true, then perhaps idealized phase spaces are not intended to represent possible states of affairs at all, but rather certain structural patterns. This attenuates the conflict with realism, since it means that distinct kinematic pairs merely represent different such patterns instead of mutually inconsistent possibilities. However, I admit that this response is somewhat sketchy: further research into the interplay between coalescence and idealization is required in order to decisively settle this question.

9 Conclusion

The story Ruetsche tells is that the ideal of pristine interpretation holds sway until one considers the UIRs that occur in QM_∞. On the contrary, I have argued that the pristine ideal struggles to account for our interpretational practices even in pre-QM_∞ scenarios, from classical and statistical mechanics to 'ordinary' quantum mechanics. Rather than an ideal, then, pristine interpretation is a simplistic—or, in Ruetsche's terms, 'extremist'—caricature. The Coalescence Approach more accurately covers the tolerance and opportunism of physics. There simply is no good reason to restrict ourselves *a priori* to one particular phase space: different facts call for different physics. These considerations ultimately support the Coalescence Approach. Since this heterodox view has so far received little attention, its details require further investigation. I have made a start by drawing a distinction between the modest and radical versions of the Coalescence Approach. As I have argued, the former is most conducive to scientific realism. The fact that modest coalescence suffices in the pre-QM_∞ contexts I have discussed as well as in the example of phase structure thereby moderates some of Ruetsche's more radical pronouncements.

However, although the ferromagnet is certainly one of Ruetsche's main cases for the Coalescence Approach, it is not her only one. It remains to be seen whether other case studies require the Coalescence Approach's more radical version. Since the main point of this paper was to draw attention to analogues of the Coalescence Approach in pre-QM_∞ theories, I will not consider these further examples in detail here. But the distinction between the modest and radical Coalescence Approach exactly affords us the possibility to ask the question which of these cases support the former, and which (if any) the latter. As I have argued, the answer to this question will have important ramifications for the status of realism and the laws of nature.

Acknowledgements

I would like to thank Adam Caulton for his helpful comments and discussions on several versions of this paper. I also thank Laura Ruetsche, Lia den Daas and several anonymous referees for additional feedback. Additionally, I am grateful to audiences at the XXI International Summer School in Philosophy of Physics in Urbino, the BSPS Annual Conference in Durham and the POP-Grunch in Oxford.

Caspar Jacobs
University of Oxford
Magdalen College
High Street
OX1 4AU
caspar.jacobs@philosophy.ox.ac.uk

References

Albert, D. Z. [2000]: *Time and Chance*, Cambridge, Mass.: Harvard University Press.

Baker, D. J. [2009]: 'Against Field Interpretations of Quantum Field Theory', *The British Journal for the Philosophy of Science*, **60**(3), pp. 585–609.

Baker, D. J. [2016]: *The Philosophy of Quantum Field Theory*, Oxford Handbooks Online. Oxford University Press.

Belot, G. [1998]: 'Understanding Electromagnetism', *The British Journal for the Philosophy of Science*, **49**(4), pp. 531–555.

Bokulich, A. [2011]: 'How Scientific Models Can Explain', *Synthese*, **180**(1), pp. 33–45.

Callender, C. [2004]: 'There Is No Puzzle About the Low Entropy Past', in C. Hitchcock (*ed.*), *Contemporary Debates in Philosophy of Science*, Blackwell, pp. 240–255.

Clifton, R. and Halvorson, H. [2001]: 'Are Rindler Quanta Real? Inequivalent Particle Concepts in Quantum Field Theory', *The British Journal for the Philosophy of Science*, **52**(3), pp. 417–470.

Earman, J. [2019]: 'The Role of Idealizations in the Aharonov–Bohm Effect', *Synthese*, **196**(5), pp. 1991–2019.

Feintzeig, B. [2016]: 'Unitary Inequivalence in Classical Systems', *Synthese*, **193**(9), pp. 2685–2705.

Feintzeig, B. [2018a]: 'Toward an Understanding of Parochial Observables', *The British Journal for the Philosophy of Science*, **69**(1), pp. 161–191.

Feintzeig, B. H. [2018b]: 'On the Choice of Algebra for Quantization', *Philosophy of Science*, **85**(1), pp. 102–125.

Friederich, S. [2013]: 'Pristinism under Pressure: Ruetsche on the Interpretation of Quantum Theories', *Erkenntnis*, **78**(5), pp. 1205–1212.

Kronz, F. M. and Lupher, T. A. [2005]: 'Unitarily Inequivalent Representations in Algebraic Quantum Theory', *International Journal of Theoretical Physics*, **44**(8), pp. 1239–1258.

Landsman, N. P. [1990]: 'C*-Algebraic Quantization and the Origin of Topological Quantum Effects', *Letters in Mathematical Physics*, **20**(1), pp. 11–18.

Lewis, D. [1973]: *Counterfactuals*, Malden, Mass: Wiley-Blackwell.

Lupher, T. A. [2008]: *The Philosophical Significance of Unitarily Inequivalent Representations in Quantum Field Theory*, Thesis.

Ruetsche, L. [2003]: 'A Matter of Degree: Putting Unitary Inequivalence to Work', *Philosophy of Science*, **70**(5), pp. 1329–1342.

Ruetsche, L. [2006]: 'Johnny's So Long at the Ferromagnet', *Philosophy of Science*, **73**(5), pp. 473–486.

Ruetsche, L. [2011]: *Interpreting Quantum Theories*, Oxford, New York: Oxford University Press.

Ruetsche, L. [2013]: 'Unitary Equivalence and Physical Equivalence', in R. Batterman (*ed.*), *The Oxford Handbook of Philosophy of Physics*, Oxford University Press.

Van Fraassen, B. C. [1991]: *Quantum Mechanics: An Empiricist View*, Oxford University Press.

Wallace, D. [2014]: ‘Review of Interpreting Quantum Theories, by Laura Ruetsche’, *The British Journal for the Philosophy of Science*, **65**(2), pp. 425–428.

Wallace, D. [2017]: ‘The Nature of the Past Hypothesis’, in K. Chamcham, J. Silk, J. D. Barrow and S. Saunders (*eds*), *The Philosophy Of Cosmology*, Cambridge: Cambridge University Press, pp. 486–499.