Review of Jonathan Bain's *CPT Invariance and the Spin-Statistics Connection*

Noel Swanson*

Jonathan Bain, ***CPT Invariance and the Spin-Statistics Connection.*** Oxford: Oxford University Press (2016), viii+194 pp., $84.95 (cloth).

The CPT and spin-statistics theorems are two important structural results concerning the foundations of relativistic quantum field theory (QFT). The first says that any well-behaved relativistic QFT must be invariant under a combined reflection symmetry that reverses the direction of time (T), flips spatial orientation (P), and conjugates all fundamental charges (C). The second establishes a connection between a particle's spin and how groups of indistinguishable particles act under permutation symmetry. Particles with integer-valued spin must obey symmetric Bose-Einstein statistics, while particles with half-integer spin must obey antisymmetric Fermi-Dirac statistics.

Both theorems help explain important physical phenomena. The CPT theorem illuminates why every fundamental particle must have an antiparticle partner with the same mass, spin, and lifetime, while the spin-statistics theorem underwrites the Pauli exclusion principle and the structure of the periodic table. On the surface, they appear to be about completely different things, but the theorems are in fact intimately linked. Foundational assumptions about the compatibility of Lorentz invariance, causality, and energy positivity are key ingredients in both results. In some proofs, CPT invariance is used to prove the spin-statistics connection. In others, the order is reversed. The lack of analogous theorems for nonrelativistic QFTs and classical field theories strongly suggests that the theorems capture constraints essential for unifying relativity and quantum mechanics.

Received March 2018.

*To contact the author, please write to: Department of Philosophy, University of Delaware, 24 Kent Way, Newark, DE 19716; e-mail: nswanson@udel.edu.

Philosophy of Science, 85 (July 2018) pp. 530–539. 0031-8248/2018/8503-0011$10.00
Copyright 2018 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/697743 Published online by Cambridge University Press

In spite of their importance, the theorems have received relatively little attention from philosophers of physics. Jonathan Bain's body of work on the subject is a notable exception, and his recent book, *CPT Invariance and the Spin-Statistics Connection*, represents an exciting addition to the philosophy of QFT literature. The book surveys and critiques four different approaches to proving the CPT and spin-statistics theorems—the textbook Lagrangian method, Weinberg's S-matrix approach, axiomatic Wightman QFT, and algebraic QFT. Ultimately, Bain argues that this plurality of approaches undermines conventional wisdom; the theorems do not explain why nature obeys CPT invariance and the spin-statistics connection. A significant obstacle facing interpreters of QFT is the lack of a unified mathematical framework for the theory, and much work in the philosophy of QFT simply picks one formulation to work with. All too often this narrowness of vision obscures important physical insight gained by looking at relations between different frameworks. Bain's book is to be commended for venturing into this underexplored territory, even if some of its more sweeping conclusions, I believe, are premature. Its careful disentangling of various proof strategies in chapter 1 and its masterful dissection of their failure in nonrelativistic theories in chapters 3 and 4 are reason enough to pick it up.

Bain has really written two books in one. The first is a detailed foundational study of the CPT and spin-statistics theorems. The second uses this investigation as a lens to explore topics in general philosophy of science including theoretical interpretation, mathematical rigor, intertheoretic relations, and scientific explanation. These two books are woven together continuously, and although the patient reader will benefit most from following both narrative threads, one could choose to focus on just one or the other and still learn a great deal.

There is a tremendous amount of thought-provoking material packed into just under 200 pages. Experts will appreciate having a compact guide to a convoluted corner of the mathematical physics literature. Newcomers will benefit from succinct discussions of a number of central topics in the foundations of QFT including the distinction between "pragmatist" and "purist" formulations of QFT, the physical motivation for various axioms (including frequently overlooked assumptions like cluster decomposition and modular covariance), and the many roles played by group representations, from classifying particle types to characterizing limiting relations between theories. Perhaps more so than any other recent philosophical monograph on QFT, *CPT Invariance and the Spin-Statistics Connection* presents the reader with a view about how all of these moving parts hang together (or in many cases, fail to do so), thus revealing exactly what is at stake in some of the more arcane foundational debates in the field.

That being said, I do not entirely agree with the synoptic view Bain sketches. While his analysis is wonderfully attuned to the subtle differences

https://doi.org/10.1086/697743 Published online by Cambridge University Press

between proofs in each framework, on balance, I worry that it overstates the divisions between these approaches and underemphasizes what they have in common. From an alternative vantage point emphasizing these commonalities, the landscape looks very different. Although it is too soon to tell which view is right, it is enough to put pressure on two of Bain's most provocative conclusions—that Lorentz invariance is not needed to prove either theorem (chap. 2) and that existing proofs do not provide explanations of CPT invariance and the spin-statistics connection (chap. 5).

Throughout the book, Bain repeatedly emphasizes how the four frameworks are mathematically and conceptually distinct: "each of these approaches can be associated with a distinct way of understanding what a relativistic QFT is about, i.e., what the basic objects of a relativistic QFT are, and what principles these basic objects are supposed to satisfy" (18). Weinberg's approach focuses on the S-matrix, which describes the dynamics of particle scattering, while algebraic QFT gives axioms characterizing the net of local observable algebras. Although both Lagrangian and Wightman QFT emphasize field operators, they assign them different properties. As a result, according to Bain, they all tell different, incompatible stories about what grounds CPT invariance and the spin-statistics connection.

A striking example of this incompatibility concerns the role of Lorentz invariance. In Lagrangian, S-matrix, and Wightman proofs, one of the primary assumptions (in slightly different forms) is that the dynamics are invariant under the group of *restricted Lorentz transformations* (i.e., rotations and boosts). In contrast, the standard algebraic proof (Guido and Longo 1995) relies on a different, apparently weaker assumption, *modular covariance*.

In algebraic QFT, as a consequence of the Reeh-Schlieder theorem, every local algebra of observables has a special pair of invariants: an antiunitary *modular conjugation* operator, J, and a one-parameter unitary *modular automorphism group*, $\{\Delta^{it}\}_{t\in\mathbb{R}}$. Modular covariance requires that the net of observable algebras must be covariant under the action of the modular automorphism groups associated with unbounded space-like wedge regions. (Such regions can be easily visualized in three dimensions as infinitely long wedges nestled into the sides of light cones.) Additionally, it requires that the associated modular unitaries, Δ^{it}, act in a geometric manner similar to wedge-preserving Lorentz boosts.[1] According to Bain, modular covariance is therefore tantamount to assuming that the net of observable algebras is covariant under a proper subset of restricted Lorentz transformations (64). As a result, he concludes that restricted Lorentz invariance is not necessary for

1. Formally, modular covariance requires that for any space-like wedge, W, and any local algebra, $\mathfrak{R}(O)$, in the vacuum representation, $\Delta_W^{it}\mathfrak{R}(O)\Delta_W^{-it} = \mathfrak{R}(\Lambda_W(t)O)$, where $\Lambda_W(t)$ is the unique one-parameter group of W-preserving Lorentz boosts.

https://doi.org/10.1086/697743 Published online by Cambridge University Press

proving the CPT and spin-statistics theorems. At the end of chapter 2 Bain tempers this conclusion somewhat, noting that Lorentz invariance is derived in the course of the algebraic proof (77), but this important caveat never receives the attention it deserves. In fact, we have good reasons to suspect that modular covariance and Lorentz invariance are more closely linked than Bain lets on.

First, if modular covariance simply amounts to the requirement that the observable net is covariant under wedge-preserving boosts, then it is not at all surprising that restricted Lorentz invariance follows. The subgroup of wedge-preserving boosts generates the restricted Lorentz group.[2] Thus, covariance under wedge-preserving boosts is not really a weaker assumption at all.

Second, the physical interpretation of modular covariance is not this direct. By itself, the condition does not entail that the wedge modular unitaries act as wedge-preserving boosts, only that they map an arbitrary local algebra onto the algebra of a suitably boosted region. Proving that they are equivalent to the corresponding boost requires background assumptions—*isotony* and *microcausality*—along with a detailed argument exploiting the special algebraic and analytic properties of modular invariants.[3] Both isotony and microcausality are standard axioms of algebraic QFT, as well as crucial ingredients in the algebraic CPT and spin-statistics theorems. The first requires that any observable localized in region O is also localized in any larger region containing O. The second enforces relativistic causality by requiring space-like separated local observables to commute. Together with these axioms, modular covariance entails the existence of a positive energy representation of the restricted Lorentz group plus the space-time translations (i.e., the *connected Poincaré group*) acting covariantly on the net. It is for these reasons that Guido and Longo (1995) characterize modular covariance as

2. The product of two noncollinear boosts is a mixture of a boost and a rotation. Consequently, the restricted Lorentz group can be generated by infinitesimal boosts in three linearly independent space-like directions. Each such boost is part of the stabilizer subgroup for some space-like wedge.

3. The modular invariants are part of the unique polar decomposition, $S = J\Delta^{1/2}$, of the Tomita conjugation, S, which implements the canonical C^*-involution when a von Neumann algebra, $\mathfrak{R}$, is concretely represented on a Hilbert space with a cyclic, separating vector, Ω. In algebraic QFT, the Reeh-Schlieder theorem guarantees that this is always the case for any local observable algebra in the vacuum representation (or more generally in any representation associated with a state analytic for the energy). The adjoint action of J maps the algebra onto its commutant, $J\mathfrak{R}J = \mathfrak{R}'$, and Ω satisfies the Kubo-Martin-Schwinger condition relative to the flow of the modular automorphism group. This entails that the vector-valued function $t \to \Delta^{it}A\Omega$, $A \in \mathfrak{R}$, has an analytic continuation into the complex strip $(-1/2, 0)$, while the function $t \to \Delta^{it}B\Omega$, $B \in \mathfrak{R}'$, has an analytic continuation into $(0, 1/2)$. For a detailed analysis of how these properties are used in the algebraic CPT theorem, see Borchers (2000).

https://doi.org/10.1086/697743 Published online by Cambridge University Press

"a way to intrinsically encode the Poincaré covariance property in the net structure" (518).

So modular covariance only indirectly tells us about space-time symmetries in conjunction with other axioms, but insofar as it does, it ensures that there is a canonical positive energy representation of the connected Poincaré group acting on the net. But it does even more than this. It also tells us that this representation is implemented by unitary operators that have unique analytic properties as modular invariants. Indeed, these extra properties are the key to showing that the representation can be extended to include an antiunitary CPT operator. From this angle, modular covariance looks much stronger than restricted Lorentz invariance. It looks more like restricted Lorentz invariance, plus translation covariance, plus energy positivity, plus additional analyticity assumptions.[4]

On top of this, Bain's discussion of the physical motivation behind modular covariance (67–69) overlooks one of the most important lines of thought—that it is not a fundamental property at all but rather can be derived from more basic algebraic assumptions. Using only translation covariance and the *spectrum condition*, Borchers (1992) proves that the modular automorphism group must act as wedge-preserving boosts, except possibly in the spatial direction along the edge of the wedge. One of the fundamental assumptions common to all four frameworks, the spectrum condition enforces energy positivity by requiring the generators of the translation subgroup to have support in the momentum-space forward light cone. Although counterexamples to modular covariance exist, none of them satisfy the conjunction of restricted Lorentz invariance, the spectrum condition, and the *split property*. The last is a technical condition ensuring the existence of normal product states across distant regions and is expected to hold in any QFT with reasonable thermodynamic behavior.[5] It is also sufficient to rule out certain counterexamples to the CPT and spin-statistics theorems constructed using infinite-component field systems (Streater 1967; Oksak and Todorov 1968) and is thus of independent interest in the

4. This idea is reinforced by considering an alternative modular assumption used in some algebraic proofs, the *condition of geometric modular action* (CGMA; Buchholz et al. 2000). CGMA is weaker than modular covariance—in conjunction with isotony and microcausality, modular covariance entails CGMA but not vice versa. For the class of generalized free QFTs, Gaier and Yngvason (2000) prove that CGMA is actually equivalent to restricted Lorentz invariance. Proving the CPT and spin-statistics theorems using CGMA, however, requires additional assumptions. An intriguing possibility is that these additional assumptions correspond to the extra analyticity properties that are built into modular covariance.

5. Formally, the split property requires that if space-time region O_1 is strictly contained in region O_2, there is a type I von Neumann algebra $\mathfrak{N}$ such that $\mathfrak{R}(O_1) \subset \mathfrak{N} \subset \mathfrak{R}(O_2)$. (In general, local algebras must be type III.) For a discussion of the physical motivation, see Haag (1996), Chap. V.5.

https://doi.org/10.1086/697743 Published online by Cambridge University Press

present debate. It remains an open question whether modular covariance can be derived from these three axioms.

To sum up, I think portions of chapter 2 need to be read with a considerable sprinkling of salt. Lorentz invariance, disguised as modular covariance, is still a central component of the algebraic proofs. Moreover, the spectrum condition, which is absent from Bain's executive summary (tables 1.1 and 1.2 on p. 28), also plays an important role. Bain argues that the spectrum condition, like Lorentz invariance, is derived rather than assumed in algebraic proofs, but if we ultimately want to derive modular covariance from more fundamental axioms, the spectrum condition is almost certainly essential. Even more importantly, in order to frame any geometric assumptions about modular invariants in the first place, these objects must exist. The physical reason why they do is given by the Reeh-Schlieder theorem, whose proof crucially relies on the spectrum condition (along with isotony and microcausality). Our discussion also brought to the fore two additional ingredients of the algebraic proofs that remain hidden in Bain's survey, the split property and analyticity (encoded by modular covariance).

These observations already serve to undermine some of the principal disanalogies Bain focuses on, but I think we can go further. In each version of the CPT and spin-statistics theorems, the following five assumptions (in some form) do most of the heavy lifting:

i) Restricted Lorentz invariance
ii) Spectrum condition
iii) Causality
iv) Finite particle multiplicity
v) Analyticity

Having just clarified the physical content of modular covariance, restricted Lorentz invariance (i) and the spectrum condition (ii) now appear in all four proof types, either explicitly or implicitly. In Wightman, Lagrangian, and algebraic proofs, causality (iii) is captured by requiring observables to commute and fields to either commute or anticommute at space-like separation. In the S-matrix approach, it is secured by Weinberg's cluster decomposition principle, which can be viewed as a kind of asymptotic version of field commutation conditions.[6]

The finite multiplicity constraint (iv) is needed to exclude the aforementioned counterexamples constructed by Streater (1967) and Oksak and

6. Clustering properties of states are dual to asymptotic abelianness conditions for observables. These were originally introduced as an asymptotic weakening of microcausality in algebraic QFT. See Bratteli and Robinson (1981), chap. 4.3, for a detailed discussion.

https://doi.org/10.1086/697743 Published online by Cambridge University Press

Todorov (1968), although it usually only appears explicitly in Wightman and algebraic proofs. In the former it is secured by requiring QFTs to be generated by finite-component Wightman fields. In the latter it is typically motivated by the split property. It is arguably a tacit assumption in both Lagrangian and S-matrix approaches, where these kinds of counterexamples do not even appear on the radar.[7]

Analyticity assumptions (v) are not well understood from a physical perspective, and yet they are essential mathematical ingredients of all known proofs of the CPT and spin-statistics theorems. In the Wightman framework, analytic properties of the *n*-point correlation functions play a central role, while in algebraic QFT, functions of the wedge modular invariants display similar analytic behavior. Lagrangian and S-matrix proofs also rely on analyticity assumptions, albeit in a less rigorous, more ad hoc manner. For example, standard formulas for dispersion relations and crossing symmetry of the S-matrix contain veiled analyticity assumptions.[8]

In spite of these structural similarities, assumptions i–v do appear in subtly different incarnations in each framework, and the resulting proofs display mathematically significant variation. Where Bain and I disagree, I think, is over the extent of this variation and whether the residual mathematical differences reflect a deeper conceptual divide. In a separate article (Swanson 2017), I defend the conceptual compatibility of Lagrangian and axiomatic frameworks. Here, I will highlight two broad reasons to be skeptical of Bain's position.

First, the mathematical assumptions we start out with in a given framework are not always intended to directly characterize its fundamental ontology. For example, field operators are not gauge invariant, so even though both Lagrangian and Wightman QFT begin with assumptions about fields, these are usually viewed as tools for constructing nets of gauge-invariant observables. Similarly, while certain structural features of the S-matrix are treated as conceptually central to particle physics in Weinberg's approach, they are still viewed as consequences of underlying gauge-invariant properties of fields (see Weinberg 2005, esp. chaps. 4.1, 5.1, 5.9, and 8.1). So even if our four frameworks disagree about the mathematical starting points, it is

7. Since these frameworks typically work with type I algebras exclusively, one might argue that the split property or the weaker distal split property is presupposed. Both are sufficient to rule out infinite-multiplet counterexamples. Moreover, as part of the physical motivation for his S-matrix approach, Weinberg (2005, xxi) cites its ability to reconcile the basic principles of relativity and quantum mechanics with a finite number of particle types. The confusion here highlights the importance of seeking rigorous proofs of foundational theorems.

8. Such analyticity assumptions were made explicit in the axioms for the S-matrix bootstrap program during the 1960s (e.g., Eden et al. 1966).

https://doi.org/10.1086/697743 Published online by Cambridge University Press

not clear that they actually disagree about what the basic objects of QFT are. Call this the *fundamentality problem*.

Second, it is important to distinguish between assumptions that are deemed essential properties of any relativistic QFT, assumptions that characterize classes of physically relevant models, and assumptions that are simply mathematically expedient. Many of the assumptions made by each framework arguably fall into these latter two categories.[9] In addition, mathematical physicists treat many assumptions as provisional, open to reinterpretation and revision. While the frameworks have different classes of mathematical models at this stage, there is also substantial overlap, and it is not clear whether these differences in modal scope reflect a disagreement about the essential features of a relativistic QFT or whether the frameworks simply characterize different types of QFTs. Call this the *scope problem*.

Together, these two problems put significant pressure on Bain's arguments concerning explanation in chapter 5. It is widely believed that the CPT and spin-statistics theorems explain why the corresponding properties are essential features of relativistic QFTs. Bain disagrees: "The existence of conceptually distinct alternative formulations of these theorems indicates that there is no unique derivation of these properties; and it also puts into question whether the principles used to derive these properties can be considered fundamental" (146). He contends that there is significant conceptual disagreement about what a relativistic QFT really is, and this in turn undermines the idea that existing proofs can provide us with explanations fitting into standard philosophical accounts of scientific explanation (deductive nomological, unificationist, structural, and causal).

In addition, all current frameworks face some version of what Bain calls the *existence problem*. One of the central interpretive claims from chapter 1 is that we can divide frameworks into "purist" and "pragmatist" varieties based on what they require of a model of QFT in order to say that it exists. Purist frameworks, like algebraic and Wightman QFT, demand existence in the form of a model of mathematically rigorous axioms. Pragmatist frameworks, like Lagrangian and S-matrix QFT, adopt "arguably weaker" conditions like Borel summability of perturbation series, renormalizability, or the existence of an ultraviolet fixed point (36–38). The problem is that we do not

9. Lorentz invariance, microcausality, and the spectrum condition are likely essential properties of anything we would be willing to call a "relativistic QFT" (although there is still some room for debate). In contrast, being the quantization of a classical Lagrangian field theory or having an S-matrix formulation are arguably not. (There exist models of the algebraic and Wightman axioms that lack one or both of these properties.) These constraints seem to pick out classes of physically well-behaved models similar to the standard model. Assumptions like additivity and Haag duality look more like provisional technical assumptions. The status of other assumptions like the split condition and modular covariance is still unresolved.

https://doi.org/10.1086/697743 Published online by Cambridge University Press

currently have models of realistic interacting theories satisfying the stronger purist existence conditions, nor do we think that any of the weaker pragmatist conditions are general enough to be plausible necessary constraints on models of QFT.

I agree that our understanding of relativistic QFT is still in flux and that the existence problem is a pressing concern, but I am inclined toward a more positive outlook than Bain in this case. Our discussion of core assumptions i–v indicates that proofs couched in different frameworks have a great deal of structure in common. Moreover, the fundamentality and scope problems cast doubt on the idea that remaining mathematical differences reflect disagreement about the fundamental objects and essential properties of relativistic QFTs.[10] On closer inspection, the conceptual divisions that supposedly undercut the explanatory credentials of the CPT and spin-statistics theorems begin to dissolve into the ether.

One of the details that tends to get lost in Bain's presentation is just how different the existence problem is for the pragmatist and the purist. The pragmatist has already succeeded in showing that the standard model exists according to Bain's list of weak existence conditions. Quantum electrodynamics, electroweak theory, and quantum chromodynamics are all renormalizable, and numerous arguments indicate that quantum chromodynamics also has an ultraviolet fixed point. The challenge for the pragmatist is that not all relativistic QFTs have these properties. So we do not yet know whether the core ideas from the pragmatist proofs apply in the general case.[11] In contrast, the purist has plausible candidates for rigorous axioms characterizing what all relativistic QFTs have in common. Her challenge is to show that these axioms have their intended scope by constructing models of four-dimensional local gauge theories like the standard model. We do not yet know whether the core ideas from purist proofs apply to the actual world.

From his vantage point, Bain sees the mathematical differences between purist and pragmatist proofs and projects a yawning conceptual crevasse separating them. As a result, the existence problem appears to preclude even

10. This undermines Bain's prima facie argument for incompatibility and shifts most of the weight onto a small handful of cases where he argues more directly (e.g., 27–29, 72–77, 147–49, and 159). Space constraints prevent me from discussing these systematically, but I think that the scope and fundamentality problems also significantly undermine these more direct arguments. For example, Bain contends that purist and pragmatist proofs disagree about whether CPT invariance entails the spin-statistics connection or vice versa. But if different frameworks characterize different subclasses of models, such logical variation comes as no great surprise and does not necessarily indicate conceptual disagreement.

11. We might also be justifiably concerned about the mathematical integrity of these pragmatist proofs, which are often lacking in rigor, but this is a separate concern (and one that Bain does not dwell on).

https://doi.org/10.1086/697743 Published online by Cambridge University Press

a provisional understanding of CPT invariance and the spin-statistics connection in realistic interacting QFTs. I have been advocating a different perspective, one emphasizing the structural and conceptual commonalities between frameworks. From this vantage point, the crevasse looks far more navigable. The fact that there is so much in common between purist proofs, which plausibly aim to capture the modal boundaries of relativistic QFT, and pragmatist proofs that cover empirically successful theories like the standard model gives us good reason to believe that a similar core of ideas will feature prominently in the eventual explanation of CPT invariance and the spin-statistics connection.

REFERENCES

Borchers, Hans Jürgen. 1992. "The CPT-Theorem in Two-Dimensional Theories of Local Observables." *Communications in Mathematical Physics* 143:315–32.

———. 2000. "On Revolutionizing Quantum Field Theory with Tomita's Modular Theory." *Journal of Mathematical Physics* 41 (6): 3604–73.

Bratteli, Ola, and Derek W. Robinson. 1981. *Operator Algebras and Quantum Statistical Mechanics*. Berlin: Springer.

Buchholz, Detlev, Olaf Dreyer, Martin Florig, and Stephen J. Summers. 2000. "Geometric Modular Action and Spacetime Symmetries." *Reviews in Mathematical Physics* 12:475–560.

Eden, Richard J., Peter V. Landshoff, David I. Olive, and John C. Polkinghorne. 1966. *The Analytic S-Matrix*. Cambridge: Cambridge University Press.

Gaier, Johanna, and Jakob Yngvason. 2000. "Geometric Modular Action, Wedge Duality, and Lorentz Covariance Are Equivalent for Generalized Free Fields." *Journal of Mathematical Physics* 41:5910–19.

Guido, Daniele, and Roberto Longo. 1995. "An Algebraic Spin and Statistics Theorem." *Communications in Mathematical Physics* 172:517–33.

Haag, Rudolf. 1996. *Local Quantum Physics*. Berlin: Springer.

Oksak, A. I., and Ivan Todorov. 1968. "Invalidity of TCP-Theorem for Infinite-Component Fields." *Communications in Mathematical Physics* 11:125–30.

Streater, Raymond F. 1967. "Local Fields with the Wrong Connection between Spin and Statistics." *Communications in Mathematical Physics* 5:88–96.

Swanson, Noel. 2017. "A Philosopher's Guide to the Foundations of Quantum Field Theory." *Philosophy Compass* 12 (5). doi:10.1111/phc3.12414.

Weinberg, Steven. 2005. *The Quantum Theory of Fields*. Vol. 1. Cambridge: Cambridge University Press.

https://doi.org/10.1086/697743 Published online by Cambridge University Press