

What Explains the Spin–Statistics Connection?

Jonathan Bain

*Dept. of Technology, Culture and Society
Tandon School of Engineering, New York University
6 Metrotech Center, Brooklyn, NY 11201*

Abstract. The spin–statistics connection plays an essential role in explanations of non-relativistic phenomena associated with both field-theoretic and non-field-theoretic systems (for instance, it explains the electronic structure of solids and the behavior of Einstein–Bose condensates and superconductors). However, it is only derivable within the context of relativistic quantum field theory (RQFT) in the form of the Spin–Statistics Theorem; and moreover, there are multiple, mutually incompatible ways of deriving it. This essay attempts to determine the sense in which the spin–statistics connection can be said to be an essential property in RQFT, and how it is that an essential property of one type of theory can figure into fundamental explanations offered by other, inherently distinct theories.

1. *Introduction*
2. *Why the Spin–Statistics Connection Needs an Explanation*
3. *The Spin–Statistics Theorem Does Not Explain the Spin–Statistics Connection*
 - 3.1. *The Deductive–Nomological Account*
 - 3.2. *The Unificationist Account*
 - 3.3. *Causal Accounts*
 - 3.4. *The Structural Account*
4. *What Would Explain the Spin–Statistics Connection?*
5. *Conclusion*

1. Introduction

The spin–statistics connection (SSC, hereafter) plays an essential role in explanations of non-relativistic phenomena associated with both field-theoretic and non-field-theoretic systems. For instance, it explains the electronic structure of solids and the behavior of Einstein–Bose condensates and superconductors. However, it is only derivable within the context of relativistic quantum field theories (RQFTs) in the form of the Spin–Statistics Theorem, and there are multiple, mutually incompatible ways of deriving it.¹ This essay attempts to determine the sense in which SSC can be said to be an essential property of RQFTs, and how it is that an essential property of one type of theory can figure into fundamental explanations offered by other, inherently distinct theories. Section 2 provides further incentive for seeking an explanation of SSC. Section 3 argues that the Spin–Statistics Theorem does not provide the sought after explanation, and Section 4 suggests what might do the job.

2. Why the Spin–Statistics Connection Needs an Explanation

SSC is a property that links the statistics a state obeys with the spin it possesses: It requires that a state that obeys Bose–Einstein statistics possesses integer spin, and a state that obeys Fermi–Dirac statistics possesses half-integer spin. Statistics can be encoded in creation and annihilation operators a , a^\dagger that act on multi-particle states in a Fock space by requiring,

¹ There is a sizable literature on attempts to derive SSC in the context of non-relativistic quantum mechanics. None of these attempts have been completely successful.

$$[a(\mathbf{p}), a^\dagger(\mathbf{p}')]_{\mp} = \delta(\mathbf{p} - \mathbf{p}') \quad (1)$$

for 3-momenta \mathbf{p} , where " \mp " indicates a commutator or anticommutator, depending on whether the particle states are bosonic (i.e., obey Bose–Einstein statistics) or fermionic (i.e., obey Fermi–Dirac statistics). Creation and annihilation operators that commute will create or annihilate multi-particle states that are symmetric under a permutation of single-particle substates, whereas creation and annihilation operators that anticommute will create or annihilate multi-particle states that are antisymmetric under such a permutation. In both cases, such multi-particle states are permutation invariant. In addition, the symmetric case allows, whereas the antisymmetric case does not allow, the presence of single-particle substates that agree on all their non-spatiotemporal properties (i.e., the antisymmetric case obeys Pauli's exclusion principle).

Statistics can also be understood in terms of field states. For relativistic fields, one requires that the field variables $\phi(x)$, $\phi^\dagger(x)$ commute or anti-commute at spacelike distances, depending on whether they're bosonic or fermionic:

$$[\phi(x), \phi^\dagger(y)]_{\mp} = 0, \text{ for spacelike } (x - y) \quad (2)$$

The non-relativistic version requires field variables to commute or anti-commute at equal times for spatial distances:

$$[\phi(\mathbf{x}, t), \phi^\dagger(\mathbf{x}', t)]_{\pm} = 0, \text{ for } (\mathbf{x} - \mathbf{x}') \neq 0 \quad (3)$$

Both versions guarantee that, when a Fock space representation is available, the creation and annihilation operators corresponding to the fields satisfy condition (1). Thus, to say that a field is bosonic (resp. fermionic), could mean either that, *by definition*, the field satisfies (2), or that, when a Fock space formulation is available, the corresponding creation/annihilation operators are associated with particle states that are bosonic (resp. fermionic).

Imposing condition (1) suggests the bearers of statistics are particles, insofar as it encodes statistics in a way that refers explicitly to the behavior of particle states under permutations. Imposing conditions (2) or (3) suggests the bearers of statistics are fields, insofar as it encodes statistics in a way that refers explicitly to the behavior of fields.

The spin possessed by a state can be encoded by identifying the state as a carrier of an appropriate representation of a spacetime symmetry group. A relativistic integer or half-integer spin state can be identified as a carrier of a true or double-valued representation of the Poincaré group. A non-relativistic integer or half-integer spin state can be identified as a carrier of a true or double-valued representation of the Galilei group.² The Poincaré group \mathcal{P} can be expressed as

² Two qualifications are in order. First, for Galilei-invariant quantum field theories, the symmetry group is actually a (central) extension of the Galilei group obtained by adding a generator corresponding to mass to the Galilei Lie algebra (Levy-Leblond 1967). Second, not all non-relativistic quantum field theories need be invariant under the Galilei group, so extended. In general, one can define a non-relativistic quantum field theory as a quantum field theory invariant under (an appropriate extension of) the symmetry group of a classical (as opposed to Lorentzian) spacetime (Bain 2010). The point is that there are more classical spacetimes than just the one associated with the Galilei group. The latter is Neo-Newtonian spacetime; others include Maxwellian spacetime, Leibnizian spacetime,

a semidirect product of the group $SO(1,3)$ of (3+1)-dim rotations with unit determinant, and the (3+1)-dim reals: $\mathcal{P} = SO(1,3) \ltimes \mathbb{R}^{1,3}$. The first term contains pure Lorentz boosts and the second term contains spatiotemporal translations. The Galilei group \mathcal{G} is a bit more complex, but can be expressed in a form that resembles \mathcal{P} ; namely, a semidirect product of a term that contains a rotation group, and a term that contains translations: $\mathcal{G} = (SO(3) \ltimes \mathbb{R}^3) \ltimes (\mathbb{R}^1 \times \mathbb{R}^3)$. The first term is a semidirect product of the 3-dim rotation group and the group \mathbb{R}^3 , which encodes velocity boosts. The second term is a direct product of real groups that encode temporal and spatial translations, respectively. Double-valued representations of \mathcal{P} and \mathcal{G} can be identified with true representations of their covering groups, obtained by replacing the rotation group, $SO(1,3)$ and $SO(3)$, respectively, with its covering group, $SL(2, \mathbb{C})$ or $SU(2)$, respectively.

Now that we've reviewed what SSC is, and how it can be represented, we should ask why does it need an explanation? Physicists point out that it has a "profound impact" (Zee 2010, pg. 121) in non-relativistic quantum mechanics, and non-relativistic quantum field theory:

From the microscopic structure of atoms to the macroscopic structure of neutron stars, a dazzling wealth of physical phenomena would be incomprehensible without this spin–statistics rule. Many elements of condensed matter physics, for instance, band structure, Fermi liquid theory, superfluidity, superconductivity, quantum Hall effect, and so on and so forth, are consequences of this rule. (Zee 2010, pg. 120.)

The world would be a different place if spin-one-half particles were not subject to Pauli's exclusion principle. In all fundamental branches of modern (natural) science, the connection between particle spins and multiparticle behavior plays a crucial role, and to date, no physical system violating it has ever been observed. (Kuckert 2007, pg. 207.)

According to these same physicists, this "profound impact" of SCC in *non-relativistic* quantum theories had to wait to be explained by *relativistic* quantum field theories:

...the explanation of the spin–statistics connection by Fierz and by Pauli in the late 1930s, and by Luders and Zumino and by Burgoyne in the late 1950s, ranks as one of the great triumphs of relativistic quantum field theory. (Zee 2010, pg. 121.)

[The Spin–Statistics theorem]... clarifies one of the great mysteries of non-relativistic quantum theory: the contrasting symmetry properties of the wavefunctions of particles of integer (boson) versus half-integer (fermionic) spin. (Duncan 2012, pg. 59.)

But not all authors are satisfied with this appeal to the Spin–Statistics theorem as an explanation of SSC:

Aristotelian spacetime, etc. (see, e.g., Earman 1989 for a bestiary). On the other hand, the non-relativistic quantum field theories of practical interest (for instance, the ones that play key roles in condensed matter physics) are typically Galilei-invariant.

The spin–statistics connection seems crucial to understanding the behavior of several physical systems for which relativistic considerations seem quite insignificant... Non-relativistic theories seem to adequately describe most of these systems and the spin-statistics connection has to be inserted 'by hand' when formulating these theories. (Shaji 2009, pg. 2.)

An explanation has been worked out by Pauli from complicated arguments of quantum field theory and relativity... we have not been able to find a way of reproducing his arguments on an elementary level... This probably means we do not have a complete understanding of the fundamental principle involved. (Feynman 1965, pg. 4-3.)

The Spin–Statistics theorem that these authors are referring to is the claim that any state of a physical system described by a relativistic quantum field theory (RQFT) must possess SSC. The concerns raised by Shaji and (famously) Feynman are concerns over the extent to which the Spin–Statistics theorem in RQFTs explains the presence of SSC in *non-relativistic* theories. This is significant, since the overwhelming majority of evidence for SSC comes from observations of physical systems best described by the latter. In the rest of this essay, I'd like to address the following questions:

- (I) Does the Spin–Statistics theorem explain SSC in RQFTs?
- (II) What explains SSC in non-relativistic theories?

The implicit assumption in Shaji and Feynman, and arguably, the received view among physicists, is that the answer to question (I) is "yes" and that there is currently no feasible answer to question (II). Sections 3 and 4 will argue against these positions. The goal of Section 3 is to argue that the answer to question (I) is "no": the Spin–Statistics theorem, by itself, does not explain SSC in RQFTs, at least under contemporary notions of scientific explanation. The goal of Section 4 is to argue that, Section 3 notwithstanding, an adequate explanation of SSC in non-relativistic theories can be constructed.

3. The Spin–Statistics Theorem Does Not Explain the Spin–Statistics Connection

I would now like to consider ways in which one might argue that the Spin–Statistics Theorem explains SSC. If it does, then we'd like to know how it does; in particular, the *type* of explanation it provides of SSC. The gist of this section is to argue that, by itself, the Spin–Statistics Theorem does not provide any of the types of explanation that philosophers of science have identified.

To set the stage, I would like to be clear on what I take to be the *explanandum*. I take SSC to be a property that a physical system may possess. As a matter of fact, it is observed that all physical systems possess this property. The question then arises, Why is this the case? The typical response in the physics literature adopts the following explanation:

The states of physical systems described by RQFTs are characterized by a set of *fundamental* properties, and SSC can be derived from this set insofar (*)

as, if the state of a physical system possesses these fundamental properties, then it must possess SSC.

In this explanation, the explanandum is a general regularity, as opposed to a particular fact. We wish to know why *all* physical systems (of the relevant type) possess SSC, as opposed to why any particular physical system possesses it. Claim (*) purports to explain this regularity by deriving it from a set of other regularities (alternatively, laws or principles) of the form "All physical systems possess fundamental property *X*". This derivation is called the Spin–Statistics theorem. At least four alternative formulations of this theorem have appeared in the physics literature, based on four distinct approaches to formulating RQFTs. These are summarized in Table 1.³ Bain (2013) has argued that these approaches are both mathematically and conceptually distinct. Not only do they differ on the principles needed to derive SSC, they also differ on what they take to be the objects that possess this property (fields, particles, or, in the case of the algebraic approach, local observables). Moreover, as will be explained in Section 3.2 below, these approaches also differ over how they treat interactions.

The task of the rest of this section is to compare Claim (*) with the standard accounts of explanation; namely, deductive–nomological, unificationist, causal, and structural.

Approach	Principles	Derived Property
Wightman	(a) Restricted Lorentz invariance (b) Spectrum condition (c) Local commutativity	SSC for field states
Algebraic	(a) Modular covariance (b) Additivity (c) Algebraic causality (d) Normal commutation relations	SSC for irreducible, restricted Poincaré-invariant DHR representations with finite statistics and masses.
Lagrangian	(a) Restricted Lorentz invariance (b) Spectrum condition (c) Causality	SSC for fermionic field states
	(a) Restricted Lorentz invariance (b) Causality	SSC for bosonic field states
Weinberg	(a) Restricted Lorentz invariance for <i>S</i> -matrix (b) Cluster decomposition for <i>S</i> -matrix	SSC for particle states

Table 1. Alternative formulations of the Spin–Statistics Theorem in RQFTs.

3.1. The Deductive–Nomological Account

A deductive–nomological (DN) explanation explains by virtue of a derivation from a set of covering laws together with a specification of antecedent conditions required to apply the laws to the given explanandum.⁴ DN explanations demonstrate how the explanandum is nomically

³ For the Wightman approach, see Streater and Wightman (1964), for the algebraic approach, see Guido and Longo (1995), for the Lagrangian approach, see Kaku (1993), and for Weinberg's approach, see Weinberg (1995).

⁴ In addition, the derivation must be in the form of a sound deductive argument formulated in a 1st-order formal language, and the premises must possess empirical content. In the following, for charity's sake, I will set aside

expected; i.e., how it follows necessarily from lawlike premises. There are at least two interrelated concerns with viewing the Spin–Statistics theorem as providing a DN explanation of SSC, both stemming from the fact that the theorem admits conceptually distinct formulations. One of these concerns has to do with the nature of the covering laws; in this case, the principles, such as *restricted Lorentz invariance*, *Cluster Decomposition*, *Modular Covariance*, etc., that underwrite the various versions of the theorem. The other concern has to do with the nature of the explanandum; in this case, the general regularity that all physical systems of the relevant type possess SSC.

An initial concern is that the notion of a covering law in a DN explanation is a matter of debate among philosophers of science. Among other things, laws have been construed as contingent or necessary relations among universals, as general regularities, or as regularities that underwrite the best (simplest, strongest, etc.) systematization of the facts. To be charitable, I will assume that the various principles (*restricted Lorentz invariance*, the *Spectrum Condition*, *Cluster Decomposition*, *Modular Covariance*, etc.) that underwrite alternative proofs of the Spin–Statistics theorem can be viewed as laws, however one chooses to describe the latter. The more pressing concern is the extent to which these principles can be considered fundamental.

The issue of fundamentality arises in the context of the second concern with viewing the Spin–Statistics theorem as a DN explanation of SSC. Again, I take the latter to be associated with a general regularity (alternatively, law or principle). The concern is that, notoriously, the DN account has problems in explaining general regularities. Hempel and Oppenheim's (1948, pg. 159) original account restricted the explanandum of a DN explanation to a particular fact due to the following problem (call it the Problem of Conjunctions): A law L_1 can be derived from the conjunction $L_1 \& L_2$ of that law with any other (true, empirical) law L_2 . This conjunction may itself be considered a law and hence can appear as a premise of a DN explanation of L_1 . But the derivation of L_1 from $L_1 \& L_2$ should not necessarily be taken to explain L_1 (take L_1 to be Kepler's laws and L_2 to be Boyle's Law). As Psillos (2007 pg. 135) notes, the problem of determining when a derivation of a law from other laws counts as a legitimate explanation of the former is the issue of what makes one law more fundamental than another: a derivation of L_1 from L should count as an explanation of L_1 whenever L counts as more fundamental than L_1 (thus the conjunction of Kepler's law and Boyle's Law doesn't explain Kepler's Law since, intuitively, the former is not more fundamental than the latter). One attempt at making this distinction, proposed initially by Friedman (1974), identifies a fundamental law with a unifying law. The next section will consider Kitcher's (1989) version of unificationism, which constitutes a distinct approach to explanation than DN. At this point, we need only understand the motivation for this approach; namely, that a law is unifying, and thus fundamental, if it belongs to a small set of laws from which a large body of claims can be derived. Under one gloss of this intuition, a principle like SSC is explained just when it can be uniquely derived from a set of fundamental (viz., unifying) principles.

The upshot of this discussion so far is that a DN explanation cannot explain SSC without facing the Problem of Conjunctions, and under one proposed solution to this problem, a principle is

concerns having to do with the extent to which theories in mathematical physics are capable of 1st-order formulation. The original description of DN is given in Hempel and Oppenheim (1948).

explained by a unique derivation of it from a set of more fundamental principles. The problem with trying to understand SSC in this context is that there is no unique set of first principles from which SSC can be derived in RQFTs. The existence of conceptually distinct alternative formulations of this theorem indicates there is no unique derivation of SSC; and it also puts into question whether the principles used to derive it can be considered fundamental. Intuitively, if SSC is to be explained by demonstrating how it is nomically expected from fundamental first principles, there should be a unique nomic story to tell about it. The Spin–Statistics theorem does not provide us with such a story.

3.2. *The Unificationist Account*

Does the Spin–Statistics theorem provide a unifying explanation of SSC? Under Kitcher's (1989) account, a unifying explanation explains by virtue of belonging to the most unifying systematization of the set K of claims currently endorsed by the scientific community.⁵ The most unifying systematization is called the explanatory store $E(K)$ over K . Thus to determine if the Spin–Statistics theorem provides a unifying explanation of SSC, we need to determine if it belongs to $E(K)$ for the relevant K .

The problem here is that there is no consensus on which approach to RQFTs should be adopted. In general, one can identify two basic approaches to formulating RQFTs. "Purist" (or rigorous) approaches like the axiomatic and algebraic formalisms attempt to identify a set of axioms and then construct models of these axioms that describe relevant field theories. These approaches face what may be called the *Problem of Empirical Import*; namely, *realistic interacting* models of the relevant axioms have yet to be constructed.⁶

An alternative type of approach might be called "pragmatist" (or heuristic). Examples include the approach developed by Weinberg, and the Lagrangian approach found in most textbooks. These pragmatist approaches face three related problems. First, in most cases, the goal of these approaches is to calculate the S -matrix, which encodes probabilities for particle scattering events. The expression for the S -matrix requires non-interacting multi-particle states at asymptotic times to be related to interacting multi-particle states at finite times; and this requires the introduction of infinitely renormalized parameters: call this the *Renormalization Problem*. A second problem is the fact that, for many *realistic interacting* theories, the power series expansion of the S -matrix contains divergent terms at high energies: this is called the *UV Problem*. Finally, for many

⁵ A systematization Σ of K consists of a subset of statements in K from which the rest of K can be derived. Σ is unifying if it maximizes scope, simplicity, and stringency. Scope is measured in terms of the number of conclusions that can be drawn from Σ . Simplicity is measured in terms of the size of Σ . Stringency is made precise in the following way. A systematization of K is a set of arguments Σ that instantiates a collection of *argument patterns*. An argument pattern consists of a schematic argument, a set of filling instructions, and a classification (the filling instructions explain how the schematic argument is to be interpreted, the classification identifies premises, conclusion, and inference rules of the schematic argument). An argument pattern is more stringent than another if the first is harder to instantiate than the second.

⁶ This should be qualified in the following ways. First, by a *realistic interacting* model, I mean a model for a 4-dim RQFT, like quantum electrodynamics or quantum chromodynamics, from which predictions have been derived and confirmed. Second, *non-interacting* models have been constructed. And finally, *unrealistic interacting* models have also been constructed; models for interacting theories in 2 and 3-dim, for instance (Rivasseau 2002, pp. 168-9).

realistic interacting QFTs, the power series expansion of the S -matrix may not even converge: call this the *Convergence Problem*.⁷

This lack of a consensus on how to formulate RQFTs makes it problematic to claim that the Spin–Statistics theorem belongs to Kitcher's explanatory store $E(K)$. Kitcher (1989, pg. 431) assumes K is consistent and deductively closed, and $E(K)$ is unique.⁸ We can restrict K to those claims associated with RQFTs in one of two ways: We can either restrict K to claims associated with just one approach to formulating RQFTs, or we can include in K claims associated with all approaches. Note that if we do the latter, we will include in K claims associated with both pragmatist and purist approaches. But then K will not be consistent. It will contain statements like "QED (as a model of the relevant purist axioms) does not exist", and "QED (as a pragmatist truncated perturbative expansion in the relevant interaction Hamiltonian) does exist". Thus if K is to be consistent, we will have to restrict it to the claims associated with just one approach to RQFTs, or at least one family of approaches, either pragmatist or purist. Suppose, then that K contains only those claims associated with pragmatist approaches. Then assumedly $E(K)$ contains pragmatist versions of the Spin–Statistics theorem, but not purist versions. Intuitively, purist versions will not even be in K . Moreover, if we put them in $E(K)$ by hand, we will make the latter larger without increasing its scope, thus decreasing its simplicity, and it will no longer be the most unifying systematization of K . Similarly, if we assume K contains only those claims associated with purist approaches to RQFTs, then pragmatist versions of the Spin–Statistics theorems will not be in $E(K)$.

This suggests that, in order to view the Spin–Statistics theorem as providing a unifying explanation of SSC, we will have to first choose between which approach to RQFTs to adopt, pragmatist or purist. This seems a bit too constraining. Indeed in practice, such a choice is typically not made. On the one hand, the practice among pragmatists suggests a combination of pragmatism and purity: For instance, one is typically told immediately after being presented with pragmatist versions of the Spin–Statistics theorem to consult purist versions for further details. On the other hand, purists certainly do not intend to reject all the claims associated with pragmatist approaches (particular claims about the values of scattering amplitudes, say, and in general, all the empirical claims associated with high energy particle physics that are derived using pragmatist techniques).

In short, if we assume that the set K of claims associated with RQFTs is both consistent and deductively closed, then the best systematization $E(K)$ of K will not be unique: there will minimally be purist and pragmatist versions. On the other hand, if we allow K to encompass all claims associated with RQFTs, both pragmatist and purist, then it will not be consistent. Note that we might expect this type of situation to arise in areas of immature science in which the

⁷ These pragmatist problems need to be qualified in the following ways: First, they are common to any approach that employs renormalized perturbation theory to derive predictions from most *realistic interacting* RQFTs. Second, some *realistic interacting* RQFTs (in particular, quantum chromodynamics) do not suffer the *UV Problem* and may not suffer the *Convergence Problem*. Finally, renormalization group techniques address the *Renormalization* and *UV Problems*, but the *Convergence Problem* typically remains.

⁸ If $E(K)$ were not unique, then the conjunction of Kepler's laws and Boyle's law might count as a legitimate explanation of Kepler's laws, although perhaps less unifying than, say, Galileo's law. Uniqueness of $E(K)$ entails, on the other hand, that the conjunction of Kepler's laws and Boyle's law is unexplanatory of Kepler's laws, which seems more appropriate.

body of claims has not yet been consistently systematized. In particular, it might be expected that unifying explanations of the phenomena associated with an immature science cannot be given. On the other hand, irrespective of how a distinction between an immature and a mature science is made, one would expect that RQFTs should count as prime examples of the latter, given their immense empirical success.

Thus, in general, since there is no consensus on how to formulate RQFTs, let alone on how to formulate the Spin–Statistics theorem, the claim that it offers a unifying explanation of SSC is suspect.

3.3. Causal Accounts

A causal explanation explains by virtue of specifying a cause. On the surface, it's not immediately clear how an explanation of SSC based on an appeal to the Spin–Statistics theorem can be interpreted as specifying causes. Note first that the intended explananda are regularities, and most discussions of causal explanation are restricted to explanations of particular events. There are, however, exceptions: Strevens (2008, pg. 220), for instance, identifies two types of causal explanation of regularities: a metaphysical causal explanation of a regularity identifies a "rich and suitably objective relation of metaphysical dependence" between the regularity and more fundamental laws, whereas a mechanistic causal explanation of a regularity identifies a mechanism responsible for the regularity.⁹ On the surface, the Spin–Statistics theorem makes no explicit reference to relations of metaphysical dependence, nor does it refer to mechanisms that might underwrite SSC.¹⁰

Another causal account of regularity explanation can be found in Woodward (2003). Woodward considers an adequate explanation to provide both causal and counterfactual information about its explanandum; the latter insofar as an explanation should exhibit a pattern of counterfactual dependence between explanans and explanandum. Woodward (2003, pg. 187) cites, as an example, an explanation for the regularity encoded in the general expression for an electric field due to a charged line source, $E = (1/2\pi\epsilon_0)(\lambda/r)$, where λ is the charge per unit length. This expression can be derived from Coulomb's law in conjunction with relevant boundary conditions. According to Woodward, the derivation provides both counterfactual and causal information.

⁹ Under this latter view, according to Strevens, "to explain a law is to have a kind of generalized understanding of the causes of instances of the law, what causalists often call an understanding of a mechanism". This view of mechanism underwrites Strevens' kairetic account of regularity explanation. Under another meaning of the term, a mechanism consists of a collection of entities and activities that are organized in such a way that they realize the regularity in question (see, e.g., Weber et al. 2013, pg. 59 and references therein).

¹⁰ Note that the way I initially formulated the explanation of CPT invariance and SSC in terms of the CPT and Spin–Statistics theorems in Claim (*) at the beginning of Section 3 might be interpreted as an instance of a metaphysical causal explanation. It takes the form of what Skow (2014, pg. 446) refers to as an "in-virtue-of" explanation. It purports to explain why physical systems described by RQFTs possess CPT invariance and SSC by citing other more fundamental properties (*restricted Lorentz invariance*, the *Spectrum Condition*, etc.) that "ground" the properties in question. Thus it claims that CPT invariance and SSC obtain in a physical system in virtue of that system possessing some set of more fundamental properties. However, whereas Skow (2014, pg. 447) dismisses in-virtue-of explanations of particular facts (events) as "obviously non-causal", Strevens (2008, pg. 220) claims that they fail as causal explanations of regularities: "...the facts about causal influence are more or less fundamental physical facts, or so I suppose, and thus are suitable stopping points for understanding... By contrast, there is no relation of dependence between laws that can be read off the physics in the same way."

Counterfactually, it shows how the expression would differ, depending on different boundary conditions. Causally, it grounds the possibilities of differing boundary conditions in causes, in the sense that the expression "...conveys information that is relevant to manipulating or controlling the field" (Woodward 2003, pg. 196). At first glance, an explanation of SSC by an appeal to the Spin–Statistics theorem might be thought of as another example. Taken literally this theorem demonstrates how the imposition of a set of constraints on the space of physically possible states of an RQFT entails that such states must possess SSC. The proofs of the theorem demonstrate how SSC depends counterfactually on the constraints; i.e., they show how, if one or more of the constraints are not imposed, then the possible states of an RQFT would not possess SSC. In this sense, they convey counterfactual information about these properties. But in what sense, if any, do they convey causal information?

Taken literally, the Spin–Statistics theorem entails that, if the state of a physical system is constrained in certain ways (i.e., if it is characterized by *restricted Lorentz invariance*, the *Spectrum Condition*, etc.), then it must be constrained in an additional way (i.e., it must be characterized by SSC). These constraints (in addition to others) then act to restrict the possible ways the state can evolve in time by means of dynamical equations of motion. A dynamical trajectory can be thought of as a path in a state space that connects an initial state with a final state by means of a dynamical map that encodes an equation of motion. This map defines what might be called a dynamical entailment relation between states. If dynamical entailment supervenes on causal dependence; i.e., if whenever two states are dynamically related, they are causally related, then a dynamical trajectory supervenes on what Lewis (1986, pg. 215) calls a causal history.¹¹ Thus, under a rather charitable understanding of dynamics and causes (under which we associate a dynamical trajectory with a causal history), it might be suggested that the Spin–Statistics theorem explains SSC by virtue of providing information about the possible causal histories of RQFT states.¹²

Apart from the amount of charity needed to associate dynamical trajectories with causal histories, this suggestion faces a problem with how the constraints associated with SSC are typically understood. In the four approaches to RQFTs in Table 1, there is an implicit distinction between two types of constraints imposed on a theory's state space, *kinematical* constraints and *dynamical* constraints. Both types jointly determine the space of physically possible states of the theory, but do so in different ways. One first identifies a space of kinematically possible states on the basis of symmetry principles and definitions of the types of physical systems one is describing (recall in the context of classical field theories that this involves, among other things, a map from a spacetime manifold to a space of field values, where the latter may have additional types of structures defined on it; in the algebraic approach, the kinematics is specified, in part, by the algebra of observables and various axioms). One then imposes a dynamics on the space of kinematically possible states. Again, this amounts to a map that takes initial states to final states, and defines the space of dynamically possible states as a subset of the space of kinematically possible states. In all the approaches to the Spin–Statistics theorem, the way the split between kinematics and dynamics is performed places superselection sectors (like those defined by

¹¹ Some advocates of the causal history approach of Lewis appear to associate dynamical laws with causal histories in this way (Skow 2014, pg. 461).

¹² This would entail that this theorem provides a causal explanation of SSC, under Lewis' causal account of explanation, if it weren't for the fact that the latter account is restricted to explanations of particular events.

statistics) on the side of kinematics; i.e., as constraints imposed on a theory's state space prior to the specification of the theory's dynamics.

Thus, in all the standard formulations, the Spin–Statistics theorem places constraints on *kinematically*, as opposed to dynamically, possible states. The theorem says that any state of a physical system described by an RQFT must possess SSC, *regardless* of what dynamics it satisfies (i.e., regardless of the theory's equations of motion). Thus even a charitable (and admittedly naive) association between dynamical trajectories and causal histories cannot underwrite an understanding of the Spin–Statistics theorem as providing causal information about SSC.¹³ Note that kinematical constraints do provide some information about dynamical trajectories, and hence (naively) causal histories, insofar as the dynamically possible states of a theory are a subset of its kinematically possible states. But, arguably, this information is causally irrelevant insofar as it is robust under variations of the dynamics.¹⁴ Thus the kinematical information (of this nature) associated with a theory doesn't tell you what would happen if the theory's dynamics had been slightly different.

A possible objection to the argument so far might run as follows: The distinction between kinematics and dynamics for a given theory is not absolute, but rather a matter of convention. Thus it should always be possible, in principle, to reinterpret a kinematical constraint as a dynamical constraint. This suggests that SSC can be interpreted as a dynamical constraint that restricts how the theory's possible states evolve in time. And this suggests, under a charitable understanding of the relation between dynamics and causes, that SSC is causal in nature. Spekkens (2014, pg. 2) argues for such a conventionality of kinematics. He claims that how one chooses to distinguish kinematics from dynamics makes no empirical difference; rather both kinematics and dynamics supervene on causal structure, and it is the latter in which explanatory power resides. As an example, Spekkens cites the Newtonian and Hamiltonian formulations of classical mechanics. These formulations posit distinct spaces of kinematically possible states (configuration space versus phase space), and distinct dynamics (second-order Euler–Lagrange equations versus first-order Hamilton equations), but agree on all empirical predictions. Thus:

It's not possible to determine which kinematics, Newtonian or Hamiltonian, is the *correct* kinematics. Nor can we determine the correct dynamics in isolation. The kinematics and dynamics of a theory can only ever be subjected to experimental trial as a pair. (Spekkens 2014, pg. 2.)

In principle then, SSC, and the constraints that entail it, can be interpreted as dynamical constraint, as long as we adjust the kinematical aspects of the theories it appears in appropriately. In fact, some authors have suggested an interpretation of the correlations associated with Bose–Einstein and Fermi–Dirac statistics as effective forces of attraction and repulsion, respectively. In particular, Bose–Einstein statistics allows a collection of bosons to all be in the same state,

¹³ Saatsi and Pexton (2013) suggest that in some cases of Woodward-style explanations of regularities, explanatory power resides in the counterfactual content alone. Insofar as this type of Woodward-style explanation does not provide causal information about its explanandum, it is not a causal explanation. The appeal to counterfactual dependence alone as explanatory will be considered in the next section under structural accounts of explanation.

¹⁴ Saatsi's (2015) discussion of a type of non-causal geometric explanation of particular events similarly puts an emphasis on kinematical aspects of the event that are robust under variation of its dynamics.

and this suggests an effective force of attraction. Fermi–Dirac statistics, on the other hand, entails the Pauli Exclusion Principle (PEP) which prohibits fermions from occupying the same state, and this suggests an effective force of repulsion.¹⁵ However, there are good reasons to reject this "real force" interpretation of quantum statistics as causal in nature; or at least as providing the sort of information one should expect of a causal explanation; namely, that of providing counterfactual information. Mullin and Blaylock (2003, pg. 1228) observe that the forces identified by the real force interpretation are supposed to be effects of the spatial part of a multi-particle system's wavefunction, and if the system possesses spin, its statistics must be encoded in a total wavefunction that combines spatial and spin degrees of freedom. Thus a total wavefunction may be symmetric even though its spatial part is antisymmetric, so long as its spin component is also antisymmetric. This may lead to examples of fermionic attractive forces and bosonic repulsive forces. Thus an explanation of PEP in terms of a repulsive force doesn't seem to convey counterfactual information of the relevant sort: If a state were not fermionic, it is not necessarily the case that it would not experience an effective repulsive force (i.e., under the real force interpretation, it is possible for bosonic states to experience an effective repulsive force).¹⁶

In general, one might agree that the kinematical/dynamical distinction is conventional, but still maintain that some constraints have a particular invariant status, insofar as they are robust (i.e., remain unchanged) under variations of the dynamics. Thus whether one chooses to call the constraints imposed by the Spin–Statistics theorem kinematical or dynamical may be a matter of convention, but in all theories in which they appear, they remain unaffected under changes in the dynamics. To the extent that such dynamical invariants do not track changes in dynamics, and hence (perhaps naively) changes in causal structure, such invariants seem independent of the latter.

3.4. *The Structural Account*

On a first gloss, a structural explanation explains by virtue of specifying mathematical structure of the relevant sort. Dorato and Feline (2011, pg. 161), for instance, claim that "...quantum theory provides a kind of *mathematical* explanation of the physical phenomena it is about. Following the available literature, we will refer to such explanations as *structural explanations*." They go on to give structural explanations of the Heisenberg uncertainty principle in terms of the mathematical structure of Fourier transformations, and the non-locality exhibited by entangled states in terms of the mathematical structure of the tensor product operation on a multi-particle Hilbert space. As examples of the "available literature", they cite the following statements of Hughes (1989) and Clifton (1998):

¹⁵ Mullin and Blaylock (2003, pg. 1224) examine a number of examples of purported fermion repulsion and bosonic attraction. The former include the virial correction to the ideal gas law, the electron degeneracy pressure of a white dwarf star, and the interaction between electrons in a diatomic hydrogen atom. An example of bosonic attraction is the behavior of trapped bosons when they condense to form a Bose–Einstein condensate.

¹⁶ Another attempt to make the concept of fermionic and bosonic effective forces respectable appears in the literature on Bohm's interpretation of quantum mechanics. Holland (1993, pg. 284) suggests that Bohmian dynamics explains PEP insofar as, while the dynamical Bohmian trajectories of fermions never cross, those of bosons can. The causal story is then supplied by appeal to the pilot wave as the force that explains these dynamical results. However, Brown, Sjoqvist and Bacciagaluppi (1999, pg. 223) point out that Bohmian dynamics only secures the fact that, if two trajectories do not coincide at a given initial time, they never will, and that if two trajectories do coincide at a given initial time, then they will coincide at all future times.

[A] structural explanation displays the elements of the models the theory uses and shows how they fit together. More picturesquely, it disassembles the black box, shows the working parts, and puts it together again. 'Brute facts' about the theory are explained by showing their connections with other facts, possibly less brutish. (Hughes 1989, pg. 198.)

We explain some feature B of the physical world by displaying a mathematical model of part of the world and demonstrating that there is a feature A of the model that corresponds to B, and is not explicit in the definition of the model. (Clifton 1998, pg. 7.)

On first glance, the Spin–Statistics theorem might be thought of as providing a structural explanation of SSC insofar as it shows how the mathematical structure associated with a set of principles, or a model of a set of axioms, demonstrates why states of a physical system must possess these properties. On second glance, however, it's not entirely clear how mathematical structure alone can explain. This concern is raised by Bueno and French (2012, pg. 99) who argue that it is not enough for mathematical structure to carry explanatory weight that it just stands in a representational relation to physical structure; rather, the physical structure so-represented must exhibit dependence relations of the relevant sort. Thus, for instance, Lange (2013, pg. 509) suggests that a "distinctively mathematical explanation" explains "by describing the framework inhabited by any possible causal relation". This framework is more fundamental than the causal relations that inhabit it insofar as it is supposed to show how the explanandum "...was inevitable to a stronger degree than could result from the causal powers bestowed by the possession of various properties" (pg. 487). Moreover, this framework is supposed to work "by constraining what there could be" (pg. 494). Let "the framework inhabited by any possible causal relation" be the space of kinematically possible states, and let causal relations be encoded in dynamics. One might then view the Spin–Statistics theorem as providing a distinctively mathematical explanation of SSC.

Alternatively, the relevant sort of dependence relations that Bueno and French request might be fleshed out counterfactually. Bokulich (2008, pg. 226) suggests dropping Woodward's requirement that an explanation provide causal information and retaining only the requirement that it convey counterfactual information: "...while I shall adopt Woodward's account of explanation as the exhibiting of a pattern of counterfactual dependence, I will not construe this dependence narrowly in terms of the possible causal manipulations of the system". This motivates the following notion of a structural explanation:

...a structural explanation can be understood as one in which the explanandum is explained by showing how the (typically mathematical) structure of the theory itself limits what sort of objects, properties, states, or behaviors are admissible within the frame-work of that theory, and then showing that the explanandum is in fact a consequence of that structure. (Bokulich 2011, pg. 40.)

Thus one might claim that the Spin–Statistics theorem provides a structural explanations of SSC by showing how the kinematically possible states of a physical system described by an RQFT are

constrained (in terms of counterfactual dependencies) to those that possess SSC.¹⁷ Moreover, Section 3.3 argued that the information that the theorem conveys is not causal (under a charitable understanding of causal information).

It thus seems clear that the Spin–Statistics theorem comes closest to providing a structural explanation of SSC. The theorem demonstrates how a set of principles constrains the kinematically possible states of a physical system described by an RQFT to possess SSC, irrespective of the dynamics such states obey. However, on closer inspection, problems arise for such a structuralist interpretation.

Note first that the conceptually distinct approaches to the Spin–Statistics theorem suggest there are distinct, competing mathematical structures that can be associated with SSC. On Bueno and French's (2012) view, it then seems strange to say there are distinct physical structures that underwrite SSC. Under Lange's (2013) view, it seems strange to say there are distinct frameworks that any possible causal relation may inhabit. This complaint assumes that the different approaches to the Spin–Statistics theorem are not just conceptually distinct, but also distinct at the level of mathematical structure. I take this to be the case: naively, taken literally, the mathematical structures associated with Lagrangian field theory, axiomatic Wightman fields, von Neumann algebras, and S -matrices are all distinct. However, I grant that an argument could be made that, for a particular subclass of RQFTs, all of these approaches supervene on a common underlying mathematical structure. In particular, there are models of noninteracting and unrealistic interacting RQFTs that do not face the *Existence Problem*. For such theories, one might attempt to identify a common underlying mathematical structure (perhaps encoded in the algebraic approach). However, the types of theories that seem to be of primary interest for explanatory accounts of SSC are realistic interacting RQFTs, and interacting NQFTs and NQM. These are the types of theories that the empirical evidence indicates possess SSC, and for these theories (at least currently) it will be hard to make the case for a common underlying mathematical structure.

The upshot of this discussion is that, for a subclass of physical systems that includes all non-interacting, and some unrealistic interacting systems, the Spin–Statistics theorem can be viewed as providing a structural explanation of SSC. For this subclass of systems, the theorem imposes constraints on kinematically possible states in such a way as to limit these states to those that possess SSC. These constraints are kinematical, insofar as they are robust under variations in the dynamics (as long as such variations remain within the subclass of systems). But one may ask, do these structural explanations underwrite an understanding of SSC? The answer, arguably, is no. In the RQFT context, the systems of interest; those that make contact with empirical tests, lie outside the subclass of systems for which the Spin–Statistics theorem provides a structural explanation. In the NQM and NQFT contexts, the systems of interest likewise lie outside this subclass. To underwrite an understanding of SSC then, we need to move beyond an appeal to the Spin–Statistics theorem.

¹⁷ This also seems to be Railton's (1980, pg. 350) notion of structural explanations as "explanations based upon laws that limit the possible states or state-evolutions of systems". According to Railton, examples of constraint laws of this kind include the Heisenberg uncertainty principle, the Pauli exclusion principle, Gibbs' phase rule, the first, second, and third laws of thermodynamics, conservation laws, and extremal principles.

4. What Would Explain the Spin–Statistics Connection?

SSC can be derived in multiple ways for non-interacting (and some unrealistic interacting) RQFTs. It cannot be derived for realistic interacting RQFTs, and it cannot be derived for interacting NQFTs and interacting theories of NQM. This is problematic, since the experimental evidence for SSC comes from interacting NQFTs and NQM. I've just argued that these considerations entail that a full understanding of SSC must go beyond an appeal to the Spin–Statistics theorem. But if not such an appeal, then *what* explains this property? It will help to first distinguish two parts to this question; namely,

- (1) Why do systems described by interacting NQFTs and NQM exhibit SSC?
- (2) Why do systems described by realistic interacting RQFTs exhibit SSC?

Towards answering question (1), I'd like to consider an example discussed by Weatherall (2011). Weatherall is concerned with explaining the equivalence between inertial mass m_i and gravitational mass m_g in the context of Newtonian gravity.¹⁸ In Weatherall's account, this equivalence is a general observational feature of the world that, while expressible in Newtonian gravity, is taken to be a brute empirical fact in that theory. On the other hand, while this equivalence is not expressible in general relativity (GR), it can be shown to arise when one considers Newtonian gravity as a limiting case of GR. Thus, according to Weatherall,

The explanatory demand is to show how, given some superseding theory, a general fact as expressed within one theory is really necessary or to be expected within the regime in which the old theory is successful... The explanatory work, then, is done by presenting the details of the relationship between the two theories. (Weatherall 2011, pg. 437.)

In Weatherall's account, a major role is played by the fact that the explanandum ($m_i = m_g$) is not expressible in the superseding theory; in Weatherall's example, this is reflected in the fact that gravitational mass is not expressible in GR. The weight of the explanation is thus carried primarily by the intertheoretic relation, as opposed to the derivation of the explanandum. In Weatherall's example, the equivalence between inertial and gravitational mass in Newtonian gravity is explained by showing how it can be derived in Newton–Cartan gravity, viewed as a limiting case of GR, and then demonstrating how Newtonian gravity can be recovered from Newton–Cartan gravity.¹⁹ The framework for this explanation is given in Figure 4.1.

¹⁸ Recall that inertial mass m_i is the constant of proportionality that appears in Newton's second law, $F = m_i a$. It measures the tendency of an object to obey Newton's first law. Gravitational mass m_g is a measure of the strength of the coupling of an object to the gravitational field, and appears in the definition of the Newtonian gravitational force $F = -m_g \partial \Phi$, where Φ is the Newtonian gravitational potential field.

¹⁹ GR is characterized by models of the form (M, g_{ab}) , where M is a differential manifold and g_{ab} is a Riemannian metric on M that satisfies the Einstein equations. Newtonian gravity is characterized by models of the form $(M, h^{ab}, t_a, \nabla_a, \rho, \phi)$, where h^{ab} and t_a are spatial and temporal metrics on M , ∇_a is a derivative operator, and ρ and ϕ are scalar fields that represent the mass density and the Newtonian gravitational potential, respectively. These objects are required to satisfy $h^{ab} t_b = \nabla_c h^{ab} = \nabla_a t_b = 0$, $R^a{}_{bcd} = 0$, and $h^{ab} \nabla_a \nabla_b \phi = 4\pi G \rho$ (Poisson equation), where $R^a{}_{bcd}$ is the curvature tensor defined by ∇_a . Newton–Cartan gravity is a geometricized version of Newtonian




Figure 4.1. Framework for an explanation of the equivalence between inertial and gravitational mass in Newtonian gravity.

In GR, the *explanandum* (the equivalence between inertial and gravitational mass) is not expressible, whereas in Newton–Cartan gravity, it is derivable, and in Newtonian gravity, it is a brute fact. Relation *A* encodes a particular limit of GR that produces Newton–Cartan gravity, and relation *B* is underwritten by a theorem due to Trautman that describes the conditions under which Newtonian gravity can be recovered from Newton–Cartan gravity.²⁰ The relation between GR and Newtonian gravity in this explanation is important: GR is the superseding theory that is assumed to provide the more accurate description of the class of phenomena associated with the explanandum. On the other hand, "...there are still regimes in which Newtonian gravitation provides a satisfactory characterization of nature" (Weatherall 2011, pg. 429). Thus, according to Weatherall, the explanation provides an answer to the question "Given that we now believe GR to have superseded Newtonian theory, then why, insofar as Newtonian theory is a limiting case of GR, are inertial and gravitational mass equal in Newtonian theory?"

An analogous question can be posed with respect to SSC in NQFTs and NQM: "Given that we now believe RQFT to have superseded NQFT and NQM, then why, insofar as NQFT and NQM are limiting cases of RQFT, does SSC hold in NQFT and NQM?" Note that, as in Weatherall's example, the empirical evidence for the explanandum appears in the theories that have been superseded (NQFT and NQM). Note, furthermore, that the explanandum (SSC) can be derived in a limited version of the superseding theory, insofar as the Spin–Statistics theorem holds only for non-interacting (and some unrealistic interacting) RQFTs. Note, finally, that the full version of the superseding theory is realistic interacting RQFTs, and, as I shall attempt to argue below, SSC cannot be expressed in the latter. Thus, in analogy with Weatherall's example, the framework for an explanation of SSC in interacting NQFTs and NQM is given in Figure 4.2.

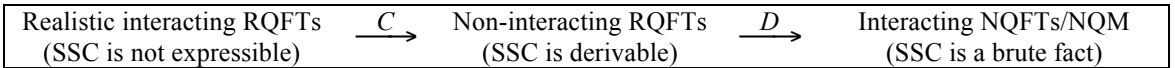


Figure 4.2. Framework for an explanation of SSC in interacting NQFTs and NQM.

gravity characterized by models of the form $(M, h^{ab}, t_a, \nabla_a, \rho)$ which satisfy $h^{ab}t_b = \nabla_c h^{ab} = \nabla_a t_b = 0$, $R^c_{abc} = R_{ab} = 4\pi G\rho t_a t_b$ (generalized Poisson equation), and $R^a_{bc} = R^c_{dab}$.

²⁰ With respect to relation *A*, the condition $R^a_{bc} = R^c_{dab}$ on the curvature tensor of a model of Newton–Cartan gravity imposes a symmetry on the connection that makes it possible to recover it as the $c \rightarrow \infty$ limit of a (general relativistic) Riemannian connection (Malament 1986). With respect to relation *B*, Trautman's recovery theorem requires an additional constraint on the curvature tensor of a model of Newton–Cartan gravity; namely, $R^{ab}_{cd} = 0$. This additional constraint is not necessary to underwrite relation *A*.

To complete this analogy, I have to accomplish two tasks. First, I have to make good on the assertion that SSC is not expressible in realistic interacting RQFTs, and second, I have to articulate the nature of the relations C and D . With respect to the former task, in non-interacting RQFTs, SSC is expressed as a relation between spin states and states that obey statistics (call the latter "statistics states"): SSC says that half-integer spin states are states that obey Fermi–Dirac statistics, and integer spin states are states that obey Bose–Einstein statistics. Recall from Section 2 that in non-interacting RQFTs, spin states are represented by carriers of true representations of the Poincaré group \mathcal{P} or its double-covering group. Statistics states are created or annihilated by field operators (or Fock space creation and annihilation operators) that satisfy conditions (1) or (2) of Section 2. One way of demonstrating that this way of expressing SSC is unavailable in a realistic interacting RQFT is to consider the implications of Haag's theorem. This theorem entails that a quantum field that satisfies the conjunction of the following conditions cannot be an interacting field (see, e.g., Earman and Fraser 2006):

- (a) the field belongs to an irreducible representation of the equal-time canonical commutation relations;
- (b) the field has a unique Euclidean-invariant vacuum state;
- (c) the field is unitarily equivalent to a free field;
- (d) the field is Poincaré-invariant.

In order for SSC to be attributed to the field, conditions (a), (b) and (d) must be upheld. Condition (a) allows statistics to be attributed to the field in the form of Section 2's conditions (1) or (2). Conditions (b) and (d) allow spin to be attributed to the field by guaranteeing that it can be represented as a carrier of a true representations of the Poincaré group \mathcal{P} or its double-covering group. This leaves condition (c). To deny it is to deny the possibility of constructing an S -matrix for the field (on pain of confronting the *Existence Problem*). Thus SSC cannot be attributed to an interacting field for which an S -matrix can be constructed. A few qualifications are in order at this point. First, Haag's theorem does not apply to non-interacting RQFTs or those unrealistic interacting RQFTs that are models of purist axioms; hence this argument against SSC expressibility in these latter theories fails. Second, what has been shown is that SSC, as expressed in a non-interacting RQFT, cannot be expressed in a realistic interacting RQFT associated with an S -matrix. This doesn't preclude the expression of SSC in realistic interacting RQFTs in a different way; but it's not that obvious what other ways there are of representing spin and statistics.²¹ The general point is that the *Existence Problem* for realistic interacting RQFTs precludes the use of the types of mathematical representations of states that underwrite a representation of SSC that non-interacting (and unrealistic interacting) RQFTs employ.

To complete the analogy between Figures 4.1 and 4.2, I now have to identify the appropriate intertheoretic relations C and D in Figure 4.2. Relation C seems easy enough: Given a realistic interacting RQFT, take the limit in which the interaction goes to zero. Relation D is a bit more complicated. Note first that, as I've argued in Section 2, the different ways of deriving SSC in

²¹ A third point could be raised to the effect that not all realistic interacting RQFTs can be associated with S -matrices, e.g., asymptotically free theories like QCD.

non-interacting RQFTs are conceptually and mathematically distinct, thus a single relation D will not be available. Rather, there will be one intertheoretic relation D for each approach to the Spin–Statistics theorem. As it turns out, this will not make any trouble.

Recall that a relation between non-interacting RQFTs and interacting Galilei-invariant QFTs (GQFT) can be induced by a speed-space contraction of the Poincaré group. A speed-space contraction transforms the Poincaré group into the Galilei group and induces a transformation of irreducible representations of the former into irreducible representations of the latter. It also induces a transformation that takes spacelike intervals into spatial intervals at equal times. This subsequently induces transformations that take the relativistic versions of the locality constraints of *Local Commutativity*, *Cluster Decomposition*, *Causality*, and *Algebraic Causality* into their non-relativistic versions. Section 3.3.2 argued that these facts underwrite the claim that a speed-space contraction defines a kinematical intertheoretic relation between RQFTs and GQFTs in each of the approaches to RQFTs and GQFTs reviewed in Sections 1.2 and 3.2. Moreover, Section 3.3.3 argued that a speed-space contraction also serves to underwrite an intertheoretic relation between *non-interacting* (and/or unrealistic interacting) RQFTs on the one hand, and *interacting* GQFTs on the other, and this makes it possible to "push down" essential properties of a non-interacting (or unrealistic interacting) RQFT to a corresponding interacting GQFT.

Thus in *non-interacting* RQFTs, SSC is an essential relation between spin states and statistics states. In (some) *interacting* GQFTs, SSC is a brute fact between spin states and statistics states. We'd like to explain this brute fact by an appeal to the derivation that grounds the essential relation in RQFTs, *and* in addition, an intertheoretic relation between non-interacting RQFTs and interacting GQFTs. The scaffolding for this explanation is represented in Figure 4.3. On the left we have an intertheoretic relation between non-interacting RQFTs and interacting GQFTs, and we'd like this relation to induce two relations in the figure on the right: one between relativistic and non-relativistic spin states, and the other between relativistic and non-relativistic statistics states. These induced relations will then allow us to "push down" the essential property of SSC from the *non-interacting* relativistic realm to the *interacting* non-relativistic realm.




Figure 4.3. SSC and relations between RQFTs and GQFTs.

A speed-space contraction of the Poincaré group induces a relation between integer and half-integer representations of the Poincaré group on the one hand, and integer and half-integer representations of the Galilei group on the other. Thus it transforms relativistic spin states into non-relativistic spin states. Moreover, by transforming spacelike intervals into spatial intervals

at equal times, it induces a relation between relativistic and non-relativistic versions of the statistics–locality connection (StLC). Thus it transforms relativistic statistics states in all four approaches to the Spin–Statistics theorem reviewed in Chapter 1 into non-relativistic statistics states. It thus provides the glue to stitch together the right hand side of the diagram in Figure 5.3. Moreover, importantly, it stitches together the left hand side of Figure 5.3, too. There are true and double-valued representations of the Galilei group that describe *interacting* GQFTs. In such cases, a speed-space contraction represents an intertheoretic relation between *non-interacting* RQFTs and *interacting* GQFTs. In such cases, we have the diagram in Figure 5.4, which provides the framework within which to construct an explanation of SSC in interacting GQFTs (where P is the Poincare group, G is the Galilei group, and StLC is the statistics–locality connection).

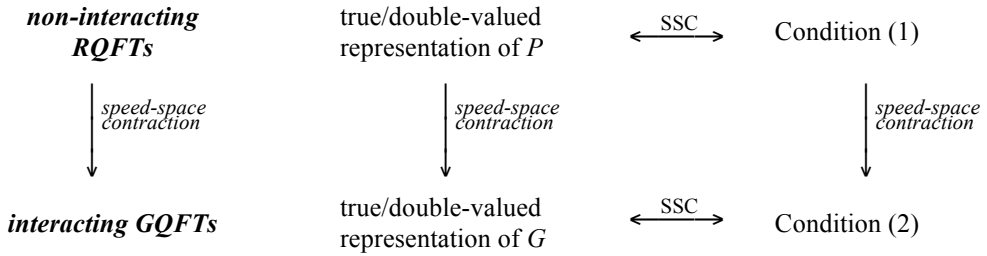


Figure 4.4. Framework for an explanation of SSC in interacting GQFTs.

Figure 4.4 shows that the general kinematical constraints shared by non-interacting RQFTs that are responsible for deriving SSC in each of the approaches to the Spin–Statistics theorem reviewed in Chapter 1 have non-relativistic counterparts under a speed–space contraction. These non-relativistic constraints explain SSC in interacting GQFTs by virtue of being the shadows of the relativistic constraints responsible for SSC in non-interacting RQFTs. Assumedly, these shadows hold equally for the finite-dimensional versions of GQFTs that take the forms of Galilei-invariant quantum mechanics (GQM).

At this point it might help to pause and take stock of the discussion so far. The first question posed at the beginning of this section was "Why do systems described by interacting NQFTs and NQM exhibit SSC?" We now have an answer to a slightly restricted version of this question, namely "Why do systems described by interacting GQFTs and GQM exhibit SSC?" This answer takes the form of a Weatherall-style explanation. This explanation explains by virtue of specifying the intertheoretic relations between realistic interacting RQFTs, non-interacting RQFTs, and interacting GQFTs and GQM. A similar answer to the original question cannot be constructed in the same way since, as Section 3.3.2 explained, the relation induced by a speed-space contraction between non-interacting RQFTs and interacting GQFTs cannot be extended to one between non-interacting RQFTs and interacting NQFTs in general; i.e., there are no similar relations between the Poincaré group and symmetry groups of classical spacetimes other than the Galilei group. This is not all that troubling insofar as the evidence for non-relativistic phenomena that exhibit SSC comes from interacting GQFTs and GQM. On the other hand, recall that in order to fully explain SSC, as well as CPT invariance, we also need an answer to

question (2) posed at the beginning of this section, "Why do systems described by realistic interacting RQFTs exhibit SSC and CPT invariance?" An answer to this latter question is more complicated. Note first that it cannot be provided by a Weatherall-style explanation, insofar as there is no intertheoretic relation between non-interacting RQFTs and realistic interacting RQFTs that could carry the needed explanatory weight. More precisely, there is no relation that transforms a non-interacting RQFT into a realistic interacting RQFT. This is the *Existence Problem* of Section 1.4.3 faced by both pragmatist and purist approaches to RQFTs. The *Existence Problem* thus entails that the constraints that the CPT and Spin–Statistics theorems impose on the kinematically possible states of non-interacting RQFTs cannot be "pushed up" to constraints imposed on the kinematically possible states of realistic interacting RQFTs, in the same way that they can be "pushed down" to constraints imposed on the kinematically possible states of realistic interacting QFTs and GQM.

Note that if there were such a relation; i.e., if there were a solution to the *Existence Problem*, then a Weatherall-style account of CPT invariance and SSC in interacting RQFTs would essentially reduce to a structural explanation. Such a relation would underwrite the claim that there is common structure shared between non-interacting and realistic interacting RQFTs, and that this structure explains CPT invariance and SSC in the latter. Assumedly, this common structure would not vary between the different approaches to formulating RQFTs, provided that a common structure underlies non-interacting RQFTs in all approaches, as was suggested in Section 5.2.4.

In light of the *Existence Problem*, then, it may be necessary to admit that currently there is no adequate explanation of CPT invariance and SSC in interacting RQFTs. This may seem like a disappointing conclusion to draw, but to maintain it is to adopt a deep understanding of the significance of CPT invariance and SSC to foundational issues in RQFTs. To realize that we currently have no explanation of these properties for interacting RQFTs is to recognize that the current status of interacting RQFTs is still very much up in the air, and that they are, in a non-trivial sense, very different types of theories than non-interacting RQFTs and interacting NQFTs and NQM.

5.4. Conclusion

What explains the spin–statistics connection (SSC)? In this essay, I've broken this question into three components. The first component asked "Does the Spin–Statistics theorem explain SSC in RQFTs?" My answer was no: This theorem does not provide an explanation of SSC in RQFTs, at least under any of the standard accounts of scientific explanation (DN, unifying, causal, or structural). My argument was based on the facts that the proof of this theorem can be formulated in mathematically and conceptually distinct ways, that the theorem does not hold for realistic interacting RQFTs, and that the theorem does not hold for non-relativistic theories (NQFTs and NQM). This conclusion generated two additional questions, which constitute the second and third components to the original question. Since much of the evidence for SSC comes from interacting Galilei-invariant QFTs (GQFTs) and Galilei-invariant QM (GQM), an explanation of SSC requires answers to the following questions:

- (1) Why do systems described by interacting GQFTs and GQM exhibit SSC?
- (2) Why do systems described by realistic interacting RQFTs exhibit SSC?

My answer to question (1) was framed in terms of Weatherall's (2011) account of explanation. According to this account, realistic interacting RQFTs presumably offer the fundamental description of the phenomena that exhibit SSC, yet SSC is not expressible, nor derivable, in these theories. On the other hand, SSC is expressible and derivable in non-interacting (and some unrealistic interacting) RQFTs. Moreover, SSC is expressible, but not derivable in interacting GQFTs and GQM. An explanation of SSC in interacting GQFTs and GQM is provided by demonstrating how it arises in a limiting process as one goes from realistic interacting RQFTs to non-interacting RQFTs (by turning off interactions), and then to interacting GQFTs and GQM, by means of a speed–space contraction of the Poincaré group. In the latter limit, in each of the mathematically and conceptually distinct approaches to the Spin–Statistics theorem in non-relativistic RQFTs, the relativistic representations of spin and statistics are transformed into non-relativistic counterparts, as are the various relativistic locality constraints (*Local Commutativity*, *Cluster Decomposition*, *Causality*, *Algebraic Causality*). A speed–space contraction transforms kinematical features of RQFTs into both kinematical and dynamical features of GQFTs. In this way, SSC as an essential property in non-relativistic RQFTs is pushed down to interacting GQFTs and GQM and thereby explained.

My answer to question (2) was less constructive: I suggested that currently there is no adequate explanation of SSC in interacting RQFTs. To recognize this is to recognize that the *Existence Problem* for both purist and pragmatist approaches to interacting RQFTs has yet to be adequately solved.

References

- Bain, J. (2013) 'CPT Invariance, the Spin-Statistics Connection, and the Ontology of Relativistic Quantum Field Theories', *Erkenntnis* 78, 797-821.
- Barcy, H. and J-M. Lévy-Leblond (1968) 'Possible Kinematics', *Journal of Mathematical Physics* 9, 1605-1614.
- Bokulich, A. (2011) 'How Scientific Models Can Explain', *Synthese* 180, 33-45.
- Brown, H. and P. Holland (2003) 'The Non-Relativistic Limits of the Maxwell and Dirac Equations: The Role of Galilean and Gauge Invariance', *Studies in History and Philosophy of Modern Physics* 34, 161-187.
- Bueno, O. and S. French (2012) 'Can Mathematics Explain Physical Phenomena?', *British Journal for Philosophy of Science* 63, 85-113.
- Clifton, R. (preprint) 'Scientific Explanation in Quantum Theory', <<http://philsci-archive.pitt.edu/archive/91/>>.
- Dorato, M. and L. Felli (2011) 'Scientific Explanation and Scientific Structuralism', in P. Bokulich and A. Bokulich (eds.) *Scientific Structuralism*, Springer, 161-176.
- Duncan, A. (2012) *The Conceptual Framework of Quantum Field Theory*, Oxford: Oxford University Press.
- Earman, J. and D. Fraser (2006) 'Haag's Theorem and its Implications for the Foundations of Quantum Field Theory', *Erkenntnis*, 64, 305–344.
- Feynman, R. (1965) *The Feynman Lectures on Physics, Vol. 3*, Reading: Addison-Wesley.
- Hempel, C. and P. Oppenheim (1948) 'Studies in the Logic of Explanation', *Philosophy of Science* 15, 135-75.
- Hughes, R. I. G. (1998) 'Bell's Theorem, Ideology, and Structural Explanation', in J. Cushing and E. McMullin (eds.) *Philosophical Consequences of Quantum Theory: Reflections on Bell's Theorem*, University of Notre Dame Press, 195-207.
- Kitcher, P. (1989) 'Explanatory Unification and the Causal Structure of the World', in P. Kitcher and W. Salmon (eds.) *Scientific Explanation*, University of Minnesota Press, 410-505.
- Kuckert, B. (2007) 'The Classical and Quantum Roots of Pauli's Spin-Statistics Relation', in W. Unruh and R. Schützhold (eds.) *Quantum Analogues: From Phase Transitions to Black Holes and Cosmology, Lecture Notes in Physics 718*, Springer, 207-228.
- Lewis, D. (1986) 'Causal Explanation', in D. Lewis (ed.) *Philosophical Papers, Vol. 2*, New York: Oxford University Press, pp. 325-47.

Railton, P. (1980) *Explaining Explanation: A Realist Account of Scientific Explanation and Understanding*. Ph.D. Dissertation, Princeton University.

Rivasseau, V. (2002) 'An Introduction to Renormalization', *Poincare Seminar 2002*, 139-177.

Shaji, A. (2009) 'Sudarshan's Non-Relativistic Approach to the Spin-Statistics Theorem', *Journal of Physics: Conference Series 196*, 012013.

Skow, B. (2014) 'Are There Non-Causal Explanations (of Particular Events)?', *British Journal for Philosophy of Science 65*, 445-467.

Spekkens, R. (preprint) 'The Paradigm of Kinematics and Dynamics Must Yield to Causal Structure', arXiv:1209.0023v1.

Weatherall, J. (2011) 'On (Some) Explanations in Physics', *Philosophy of Science 78*, 421-447.

Zee, A. (2010) *Quantum Field Theory in a Nutshell*, Princeton: Princeton University Press.