

# What good is Haag’s no-go theorem? What axiomatic methods can teach us about particle physics

Adam Koberinski\* †

## Abstract

Haag’s theorem is a no-go theorem for the interaction picture in relativistic quantum field theory. However, the interaction picture is still widely used in conventional perturbative calculations. But how exactly is the no-go theorem thereby avoided, and what do these formal results tell us about the physical systems we study, if anything? I argue that the value of axiomatic quantum field theory for modelling particle physics systems lies in understanding the structural relationships between certain features of a quantum field description. For Haag’s theorem, we learn that unitary inequivalence is related to the infinite spacetime volume idealization. We can evade Haag’s theorem using more realistic idealizations, or we can ignore it for all practical purposes when we know the root cause. Various solution strategies from axiomatic quantum field theory make this conclusion clear, where conventional Lagrangian approaches are too conceptually coarse to pinpoint the infinite spacetime idealization. However, the two approaches are consistent, indicating the value of axiomatic methods for quantum field theory.

## 1 Introduction: No-go theorems and foundations of physics

Haag’s theorem is a form of no-go theorem in quantum field theory (QFT): the standard formulation due to Haag (1955), Hall and Wightman (1957) assumes the Wightman axioms for QFT, and one major conclusion is that no free QFT can be unitarily equivalent to a theory with interactions (Earman and Fraser 2006; Freeborn, Gilton, and Mitsch 2022). This historically posed a problem for theorists using perturbative methods to generate predictions from QFT: the interaction picture assumes that the free theory is related to the full theory by a small perturbing interaction, and that the free and interacting theories are related by a unitary transformation. Unitary equivalence is usually required for different formulations of a quantum theory to ensure that probabilities of outcomes are preserved; this establishes a clear form of theoretical equivalence between the two descriptions. Insofar as the conventional, Lagrangian QFTs (LQFTs) are thought to satisfy the Wightman axioms, the theorem states that perturbative calculational techniques are mathematically inconsistent. Generalizations of the theorem have been proven for different rigorous formalisms besides the Wightman axiomatic system, suggesting that the result is robust in many plausible axiomatic QFT (AQFT) formalisms. Nevertheless, physicists have used—and continue to use—the interaction picture to perturbatively calculate scattering amplitudes, which provide the basis for the most precisely confirmed theoretical predictions in particle physics. This tension raises two sets of interesting questions: first, how is it that an apparently inconsistent formalism is used to such a high degree of success in science? Second, what good is a no-go theorem, if physicists can continue to ‘go’ anyway? The answer to the latter is informed by the answers to the former, and these questions further serve to inform how philosophers should understand and interpret physical theories and the approximations and idealizations they employ. I argue here that theoretical frameworks

---

\*Rotman Institute of Philosophy, University of Western Ontario, London, ON, Canada

†akoberin@uwo.ca

are analogous to modelling frameworks, and highlight the ways in which inequivalent frameworks can be used to model systems in physics. This sort of view of putatively fundamental physics as more analogous to applied sciences may come as a surprise, though the similarities are clear and illuminating for understanding physics practice, as well as foundational interpretations of physical systems in particle physics.

Freeborn, Gilton, and Mitsch (2022) (FGM) provide a survey of responses to Haag’s theorem in the literature, and show that many of the responses depend crucially on background assumptions about the best framework for particle physics, and perhaps the best way forward to go beyond the standard model of particle physics. Two of these camps are characterized in terms of a background adherence to AQFT over LQFT or vice versa for foundational work in this area. Applying the modelling frameworks view to particle physics, this encourages one to instead see LQFT and AQFT as two different, complementary modelling frameworks.<sup>1</sup> I argue that these two frameworks can work in tandem at the foundations of particle physics, to provide complementary perspectives on important physical aspects of the systems we study. To situate this paper within the FGM scheme, I argue that foundational work is best conceived of as using all resources at one’s disposal to further our understanding of the relationship between our best models and the world. Some aspects of a given model are representational, while others are purely formal, and still others are approximations or idealizations. Using competing modelling frameworks can help one better make these distinctions, and separate the true physical content from the formal, idealizations, and approximations.

With this understanding of the relationship between AQFT and LQFT in hand, I turn attention to Haag’s theorem (Sec. 3), and argue that the understanding of its implications from LQFT is incomplete (Sec. 3.1). By turning attention to AQFT (Sec. 3.2), we can zero in on the interpretive implications of Haag’s theorem: it represents a mathematical pathology that arises as a result of an infinite spacetime volume idealization for the quantum field description. The clever workarounds in AQFT also provide interesting hints at future physics.

## 2 Modelling frameworks in particle physics

The philosophical attitudes regarding the relative merits of the LQFT and AQFT formalisms have largely been shaped by the debate between Doreen Fraser and David Wallace from 2009-2011 (Fraser 2009; Wallace 2011).<sup>2</sup> The debate, focusing on the proper formalism for philosophical interpretation, pitted AQFT and LQFT against each other as rival research programs. On the one hand, LQFT is the basis of the empirically successful practices in particle physics, though its formulation is not amenable to standard methods of theory interpretation in philosophy of physics. On the other, AQFT offers the conceptual clarity lacking in LQFT, though there is good reason to expect that models of interactions needed for particle physics cannot be formulated in current axiomatic or algebraic frameworks (see, e.g., Aizenman and Duminil-Copin 2021). If we take an either-or approach, and our aim is to do philosophy of physical theories that apply to the world, then it seems that LQFT wins out. One can see this influence in the shift in philosophy of QFT toward engaging with LQFT and reimagining the interpretive project for philosophy of physics.

However, if we want to make sense of physics practice, we must grapple with the fact that AQFT methods are used to prove theorems thought to hold for the systems under study in LQFT. Things like the spin-statistics theorem, the CPT theorem (Streater and Wightman 1964), and the Reeh-Schlieder theorem (Reeh and Schlieder 1961) are all proven using various AQFT methods, and all

---

<sup>1</sup>I intentionally paint with a broad brush here: AQFT is a large, heterogeneous collection of different axiomatic frameworks, some utilizing algebraic formulations of QFT, and LQFT is the stand-in for the collection of theories and methods used by the majority of physicists to generate predictions far from the foundations of the field.

<sup>2</sup>In personal communication, the two spokespeople in the debate have both come to share something of a more reconciliatory attitude towards the two formalisms. However, due to the influence of the papers framing the debate as that between rival research programs, I will take some time to engage with and argue against this framing. Koberinski (2016) and Swanson (2017) have previously argued that the two should not be viewed as rivals.

are thought to be informative of the nature of QFTs. Haag’s theorem is another example that putatively bears on results for QFT generally. If these are really incompatible competing formalisms vying to describe the same systems, then how can we make sense of the widespread use of AQFT for understanding some properties of the systems we think are best described by LQFT? The answer lies in rejecting the characterization of the contrast as that between rivals, and reframing in terms of complementary modelling frameworks for complex systems, as in Figure 1.

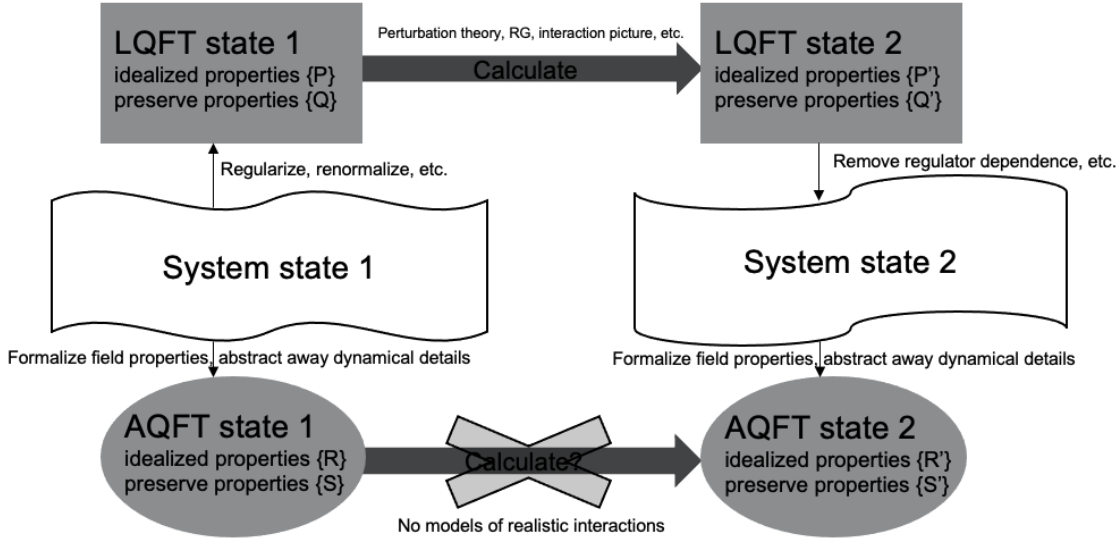


Figure 1: Schematic for thinking of LQFT and AQFT as complementary modelling frameworks, each partially describing the quantum field systems in the world. “System state x” depicts the actual system in the world, modelled variously using AQFT or LQFT. In this simple example, the system transitions from state 1 to state 2. In each of these modelling frameworks, different idealizations, approximations, and abstractions are used to represent the systems in different mathematical formalisms. While LQFT is fit for calculating dynamical evolution and modelling interactions, AQFT is only able to capture overall structural features.

It is not common to think of theoretical frameworks as modelling frameworks in the philosophy of physics. Typically, one thinks of theories as the appropriate ground for metaphysical interpretation, and that incompatible theories cannot both be true or adequate representations without explicit limiting relationships between the theories establishing something like reduction or some form of equivalence. In cases where such relationships can be established, therefore, the analogy is unhelpful. But there are important cases where these relationships cannot be established, or only partial relationships exist, while we nevertheless want to say that each theory provides partial, complementary representational capacities. Different idealizations, approximations, and abstractions are used in each framework, highlighting and preserving certain physical features while neglecting others. Each of these partial representations contributes something different to our overall understanding of the target domain. These are common in physics when thinking of complex systems with many degrees of freedom, where representations at different scales are important (Batterman 2021), or in other applied sciences where a fundamental, best representation is implausible or impossible, as in modelling weather and climate patterns.

Familiar from these other areas of science, I claim that we should think about LQFT and AQFT as distinct modelling frameworks, each using different mathematical formalisms to accurately capture certain features of the systems under study, and idealizing, approximating, and abstracting other

features. Distinct modelling frameworks can be used to model the same system, even when there is apparent incompatibility or inconsistency between them. Different idealizations might capture some relevant dynamical or structural features, while distorting others, while different approximations might distort features we take to hold in actual systems. In particular, the LQFT formalism is far more powerful for modelling interactions, while AQFT can only model abstract structural features of fields. The flexibility of LQFT is one thing that makes it a powerful formalism for physics, and allows for a wide range of physical interpretation. This flexibility makes interpretation challenging, since many complex manipulations are required to calculate or extract predictions. It is not always clear whether one ought to take these manipulations as indicating some physically significant feature of the formalism, or whether they are merely formal manipulations needed to calculate or approximate. AQFT can serve as a powerful complementary guide to physical interpretation in cases where the flexibility of the LQFT formalism hinders conceptual clarity. For capturing abstract structural features of systems, AQFT is advantageous; the logical implications of certain features (e.g., Poincaré covariance) for others (e.g., necessity of unitarily inequivalent representations) are much clearer to draw in AQFT. We should therefore use AQFT as a secondary modelling framework to aid in physical interpretation of real physical systems. One advantage of such an approach is that, when physical interpretation from both frameworks converge, we have stronger evidence that the interpretation is (approximately) true of the actual systems. In this way AQFT can either correct, or strengthen and reinforce our foundational understanding from LQFT.

Careful attention to the actual practice of physics motivates a view of the relationship between LQFT and AQFT as one of complementary modelling frameworks for relativistic quantum field systems. When engaging in foundational work, including physical interpretation, one should use both frameworks to provide a convergent understanding of the physics of these systems. The next section applies this general lesson to the specific case of Haag’s theorem, to help understand its physical implications for our understanding of relativistic QFTs. This goes some way to answering the more general question motivating this paper: what good is a no-go theorem for the foundations of particle physics?

### 3 Evading Haag’s theorem

Haag’s theorem is an interesting test case for this view of the relationship between LQFT, AQFT and physical systems in particle physics. First suggested by Haag (1955), and proven by Hall and Wightman (1957), it employs the Wightman axiomatic system to show that no free QFT can be unitarily related to an interacting QFT. This is a problem insofar as perturbative methods using the interaction picture are thought to obey the Wightman axioms or some analogue, and the interaction picture relies on unitary equivalence between free and interacting Hamiltonians. Even the setup for the importance of this no-go theorem presupposes some sort of relationship between AQFT and LQFT; the rival research program approach would dismiss Haag’s theorem as not relevant to the formalism necessary for successful physics practice.

The Haag-Hall-Wightman theorem takes the form of a *reductio ad absurdum*. Since the conclusion is unacceptable, one must reject one or more of the premises. This could be either one or more Wightman axioms, or the assumption of unitary equivalence needed to set up the standard interaction picture. The rejected premises and resulting physical interpretation will vary depending on the formalism one adopts. I will examine the ways in which the theorem is evaded in LQFT, as well as candidate modifications to the Wightman axioms which lead to its evasion in AQFT. I argue that, while the physical significance of the evasion is unclear in LQFT, supplementing this with options from AQFT leads to a deeper understanding of how different modelling assumptions remove the pathology, and suggests an interesting convergent physical interpretation. In particular, we learn that unitary inequivalence is due to an idealization of infinite spacetime volume.

I offer here only a sketch of the Haag-Hall-Wightman theorem, split into two parts. The first es-

establishes the unitary equivalence of the vacuum states, while the second then establishes the equality of the first four  $n$ -point functions of two theories. For the case where one theory is a free theory, this is sufficient to establish full equivalence. For more engaged discussion with the details of various forms of the theorem, see Earman and Fraser (2006) and Duncan (2012).

*Haag-Hall-Wightman theorem part I*

Start with two chargeless scalar fields  $\phi_i, \phi_2$  satisfying the equal time canonical commutation relations,

$$[\phi_i(\mathbf{x}, t), \pi_j(\mathbf{y}, t)] = i\delta(\mathbf{x} - \mathbf{y})\delta_{ij} \tag{1}$$

$$[\phi_i(\mathbf{x}, t), \phi_j(\mathbf{y}, t)] = [\pi_i(\mathbf{x}, t), \pi_j(\mathbf{y}, t)] = 0, \tag{2}$$

where  $\pi_i$  is the conjugate to  $\phi_i$ . Each field carries a unitary representation of the Euclidean group of spatial translations  $\mathbf{a}$  and rotations  $\mathbf{R}$ ,  $U_j(\mathbf{a}, \mathbf{R})$ . Suppose the two fields are related by a unitary transformation at some time  $t_0$ ,  $V(t_0)$ . This implies that the representations of the Euclidean group for each field are related by this same unitary  $V(t_0)$ . If each field representation has a unique “vacuum” state  $|0_j\rangle$  invariant under the Euclidean transformations,

$$U_j(\mathbf{a}, \mathbf{R}) |0_j\rangle = |0_j\rangle, \tag{3}$$

then the two vacua are unitarily equivalent.

*Haag-Hall-Wightman theorem part II*

Consider the fields from part I. Suppose they also each bear unitary representations of the Poincaré group of transformations, under which the states  $|0_j\rangle$  are also invariant, and that there are no negative energy states of the field. Then the first four  $n$ -point functions of each field are equivalent:

$$\langle 0_1 | \phi_1(x_1)\phi_1(x_2)\phi_1(x_3)\phi_1(x_4) |0_1\rangle = \langle 0_2 | \phi_2(x_1)\phi_2(x_2)\phi_2(x_3)\phi_2(x_4) |0_2\rangle. \tag{4}$$

In the case where one of  $\phi_j$  is a free field, there is a theorem from Jost (1961) that shows that any field whose two-point function coincides with the two-point function of a free field is itself a free field. This implies that the other field cannot be interacting. The conclusion is that a free and interacting fields satisfying the conditions stated (all falling under the Wightman axioms) cannot be related by a unitary transformation; if one field is free then they both are. Generalizations of Haag’s theorem exist using different axiomatizations or algebraic methods, suggesting that this is not simply a problem with the Wightman axioms.

### 3.1 . . . in LQFT

Haag’s theorem seems to pose a puzzle for using the interaction picture. Given that LQFT describes quantum fields, and that various AQFT realizations encode formal properties of those fields, *how can the interaction picture, as used by physicists, be mathematically consistent in light of Haag’s theorem?* Duncan (2012) provides a clear, detailed account of how calculations in the interaction picture are actually carried out in LQFT, and pinpoints the evasion in the necessity of regularizing the path-integral to obtain finite results using the interaction picture. Regularization is usually carried out by imposing a high-energy (UV) and low-energy (IR) cutoff on the path integral, effectively placing the LQFT on a lattice and making it a theory with finitely many degrees of freedom. In doing so, one breaks the Poincaré covariance of the theory, which was a necessary step in the schematic proof above, and renders the number of degrees of freedom finite, which also violates the assumptions of the theorem. Haag’s theorem no longer applies, and the interaction picture is well-defined. After calculating, we remove the regulators in a way that preserves the predictions and restores Poincaré covariance; when we can do so successfully, the theory is said to be renormalizable. Other regularization schemes retain the continuum, but serve to tame UV and IR divergences simultaneously.

Miller (2018) argues that regularization rescues the empirical adequacy of LQFT from the threat of mathematical inconsistency. Indeed, the solution offered by Duncan and Miller adequately addresses the question posed at the start of this section. But what are we to make of the procedures of regularization and renormalization themselves? Do they tell us something of physical salience about QFT systems? According to Miller, “[t]he best available explanation of this fact is that the observables that get compared to experiment are insensitive to the removal of the [IR] cut-off” (p. 815). So, in LQFT, we construct a continuum theory that we think is a good representation of the systems under study, render it finite to calculate, then take a continuum limit again to compare the predictions to experiment. LQFT is famous for its operationalist origins, and the difficulty of distinguishing physically salient bits of the formalism from mere calculational devices. Where do the regulators fall on this scale? Is there a relevant physical difference between the UV and IR regulators, as suggested in the quote from Miller?

There is further debate about the exact way that LQFT evades Haag’s theorem, that involves closer attention to the way that continuum limits are taken in this process. Fraser (2006) argues that the continuum limits amount to the introduction of an inconsistent mathematical representation, while Klaczynski (2016) argues that the limits introduce non-unitarity, and that it is the non-unitarity doing the work of evading Haag’s theorem. One major drawback to the LQFT formalism is that these questions do not have unambiguous answers, meaning that interpretive work is open-ended. Essentially, regularization and renormalization provide a shotgun approach to solving technical and conceptual issues for LQFT; many changes are introduced all at once, and so the logical dependence of any resulting formalism on any one specific change is obscured. I present some suggestions for interpreting the role of regulators in evading Haag’s theorem here, then supplement these suggestions with compatible interpretations from AQFT in the next section. This serves to underscore the point that the two modelling frameworks can work in complementary ways.

One option is to read the necessity of some form of regulators—both UV and IR—as directly indicative of something about the ontology of systems described by LQFT. Such a view is motivated by the idea that we should endow necessary aspects of the formalism of a successful physical theory with direct physical significance. One could then infer directly that systems described by LQFT are systems with a finite number of degrees of freedom—and therefore not really quantum *field* systems at all! But this move is too quick for a few reasons. First, though the presence of *some* sort of regulator is necessary for carrying out calculations, that regulator need not be a sharp cutoff. Even softer cutoffs, like those that exponentially dampen UV or IR modes, make this sort of interpretation much harder. With soft cutoffs, all of the continuum degrees of freedom are there, though many are suppressed. Other regulators introduce fictitious particles or change the dimension of spacetime, and these become even harder to make sense of in terms of strictly finite degrees of freedom. Second, regardless of the choice of regulator, the empirically relevant predictions extracted from LQFT are at least highly insensitive to the value of the regulator, and usually the regulator is removed before comparison with experiment. So this more direct ontological significance seems too hasty.

But a more sophisticated version of the above is to endow the regulators with physical significance in a looser, epistemic way. The necessity of regulators in LQFT blocks us from taking them to be even candidate descriptions of the world valid to all scales, while considerations external to the theory lead us to expect that even as candidate descriptions, they would fail of the actual world anyway. Thus we should take the necessity of some form of regulator to be a signal of the breakdown of our LQFT description in terms of the chosen particle and fields at some energy scale. A particular choice of regulator—such as a cutoff or dimensional regulator—is then a simple way of parameterizing our ignorance that makes the model tractable. Cutoff regulators are types of idealizations, ones that break the approximate field description and the exact symmetries of the theory; we restore the field description and the symmetries after calculating, but we should also be aware that these only hold approximately and relative to a scale (Wallace 2021). Dimensional regulators are more clearly formal idealizations used to make integrals converge, but have the advantage of preserving the symmetry structure of the theory. We can justify the use of the field description in terms of renormalization

group analysis, which indicates the robustness of low-energy field descriptions in LQFT (Koberinski and Fraser 2023). This is a much more sophisticated and apt interpretation of the LQFT formalism, but it is still somewhat unhelpful for pinpointing the problem the Haag’s theorem raises. What does our ignorance of high-energy physics imply about interpreting QFT or the interaction picture? If these are approximations or idealizations, how can they imply anything about physical possibility?

Take the solution as stated by Duncan (2012), for example. He claims that,

the proper response to Haag’s theorem is simply a frank admission that the same regularizations needed to make proper mathematical sense of the dynamics of an interacting field theory at each stage of a perturbation calculation will do double duty in restoring the applicability of the interaction picture at intermediate stages of the calculation. (p. 370)

This tells us that the presence of regulators violates the assumption of Poincaré covariance of the theory, thereby evading Haag’s theorem. Additionally, the imposition of UV and IR regulators reduce the number of degrees of freedom in the theory, down to a large but finite number. In this case, the Stone-von Neumann theorem applies, assuring us that different Hilbert space representations of the theory are unitarily equivalent. So the introduction of regulators serves as a blunt instrument, making many changes to the theory and solving numerous mathematical pathologies. But this answer doesn’t fully pinpoint what specific features of our modelling framework are to blame for Haag’s theorem, and the exact role that regulators play in violating these assumptions. First, the presence of *either* a UV or IR regulator on its own would be sufficient to break Poincaré covariance, but generically will not guarantee that perturbative calculations are well-defined or make the theory finite. So we don’t learn from this whether Haag’s theorem is related to UV divergences, IR divergences, or some combination of both. We don’t even learn whether breaking Poincaré covariance alone is sufficient, or whether it is also necessary to make the theory finite. Second, alternative regularization schemes exist that retain the symmetries of the theory, including some appropriate generalization of Poincaré covariance. In particular, dimensional regularization is widely used as a regularization scheme for precisely this reason. Dimensional regularization also ensures that perturbative calculations are well-defined, but does so in a way that respects the symmetry structure of the theory. Since physically meaningful quantities are supposed to be insensitive to the choice of regularization scheme, we have no ground for privileging cutoffs over dimensional regularization. Therefore, the evasion of Haag’s theorem can’t be only explicable for LQFT with cutoffs.

Despite all of these confusions, the evasion via cutoff regulators rendering the theory finite is suggestive of interpretive moves that one *could* make for understanding particle physics systems via LQFT. The best candidate answer we get is that imposing UV and IR cutoffs is sufficient for evading Haag’s theorem. But the formalism blends together many ingredients in ways that make a clear separation of dependencies rather challenging. We don’t get a clear understanding of what changes are simply mathematically necessary approximations without direct physical content, and what changes are physically significant. Luckily, Haag’s theorem was originally formulated in the AQFT framework, where dependencies are much clearer.

### 3.2 ... in AQFT

The Haag-Hall-Wightman proof of Haag’s theorem employs the Wightman axioms, showing the impossibility of modelling a free field and interacting field within the same Hilbert space. But like any no-go theorem, one can respond by rejecting at least one of the premises needed to prove it. We saw for LQFT that standard regularization schemes break Poincaré covariance, and techniques required to calculate generally lead to changes to the mathematics of the theory in other ways that would violate the Wightman axioms. But the evasions in LQFT don’t tell us much about how to interpret the physical significance of the theorem, or what LQFT calculational techniques imply for our understanding of particle physics systems in the world. AQFT methods provide some insight

into the structural features of quantum fields; insofar as the systems we describe are successfully represented as quantum fields, we expect many of these structural features to hold. Despite Haag’s theorem seemingly forbidding a certain representation of *dynamical* interactions, it actually implies something *structural*: a Hilbert space with unique vacuum state with respect to one Hamiltonian cannot also contain a different unique vacuum state with respect to a different Hamiltonian. Since this does not directly relate to parts of AQFT we know to be inappropriate for describing real world systems, we can use the conceptual clarity of AQFT to understand how generalizations of the Wightman axioms can avoid Haag’s theorem. Different solutions will clearly violate one or more of the assumptions that go into the Haag-Hall-Wightman proof. By analyzing a set of different solutions, the common ground they share provides evidence for what assumptions are individually necessary and jointly sufficient to prove or evade Haag’s theorem. In agreement with the conclusion of Miller (2018), we see that AQFT suggests that unitary inequivalence is related to the IR behaviour of a QFT. Specifically, we can evade Haag’s theorem by giving up the infinite spacetime volume idealization commonly employed in both AQFT and LQFT.

Earman and Fraser (2006) provide several options for relaxing or modifying the Wightman axioms that serve to evade Haag’s theorem. I will discuss a few of these options, with the aim of combining the insights here from those suggested by LQFT. The more recent criticisms of their account are correct, but have been overblown; while, e.g., Miller (2018) is correct to point out that they leave open the question of understanding LQFT approaches to evading Haag’s theorem, their conclusions and framing are compatible with a conciliatory understanding of the relationship between AQFT and LQFT. Using the conclusions from Sec. 2, we can make this more precise. When we use both modelling frameworks in a complementary way, we can get deeper physical insight into particle physics systems and their properties.

The first option bites the bullet and accepts that the interaction picture is untenable for QFT. Instead, Haag (1958) and Ruelle (1962) developed a new scattering theory, where one is always within the Hilbert space for the interacting theory only. Instead of a unitary transformation from free fields “at infinity” to interacting fields in the scattering region, one defines a weaker notion of a surrogate state, and identifies regions of the interacting Hilbert space with the span of surrogate free states. Haag’s theorem is not an issue, because the notion of equivalence here is not unitary equivalence. We know that the interaction picture is itself an idealization: the modelling assumption that we can turn the interaction off at infinity is blatantly unphysical. We use free states to have a well-defined notion of particle within the theory, but these are also approximate. Haag-Ruelle scattering theory is a construction that proves that a more careful modelling framework—including a different notion of approximately free states—makes a revised scattering picture mathematically well-behaved. Though unitary equivalence seems like a good option for treating these approximations, a weaker notion of equivalence is also justifiable. One major drawback here is that Haag-Ruelle scattering theory is rather complicated, and still only deals with asymptotic states. Thus we can retain an infinite volume idealization if we weaken the notion of equivalence required. But we gain insight into the places where modelling idealizations start to cause problems for the well-definedness of the theory.

A second option, used explicitly for constructive models within the AQFT framework, is similar in spirit to Haag-Ruelle scattering theory. Instead of requiring a global notion of unitary equivalence between free and interacting theories, one can establish a local unitary equivalence within some bounded region of spacetime (Reed and Simon 1975). This ensures that some set of local algebras for each theory are unitarily equivalent, while the full nets remain unitarily inequivalent. For toy theories without UV divergences, this local notion of unitary equivalence suffices to construct an interaction picture and evade Haag’s theorem. Since we only ever model systems of finite extent, demanding global unitary equivalence while idealizing to infinite spacetime volume might be seen as layering too many unphysical idealizations on top of one another. As we saw from Haag-Ruelle scattering, we can retain the infinite spacetime volume idealization if we reject unitary equivalence. Here we see that we can retain unitary equivalence if we reject infinite volume. We have an example of a solution that continues to treat QFT systems as continuous, but limits the extent of systems



to local regions. The lesson here is compatible with the tentative conclusions drawn from LQFT: Haag’s theorem is directly related to IR and/or UV behaviour of the idealized theory. What we learn to sharpen that suggestion from LQFT is that Haag’s theorem can be evaded by only placing IR restrictions; the UV behaviour of the theory is immaterial to the presence or absence of unitary inequivalence. Insofar as the conclusions are compatible, we should expect that introducing an IR regulator, effectively removing the infinite volume idealization in LQFT would also suffice to evade Haag’s theorem. So the lessons from AQFT here reinforce and sharpen the hints from LQFT.

One final option from generalizations of AQFT to curved spacetime settings is worth mentioning here. For spacetimes lacking some of the characteristic symmetries of Minkowski spacetime, constructing well-defined local AQFTs poses new challenges. Hollands and Wald (2010) provide an axiomatic formulation of QFT on curved spacetime, and argue that one major change is a restriction to local regions of spacetime. Since generically curved spacetimes have little global symmetry structure, this restriction to local regions is physically motivated by consistency with general relativity. They formulate the axioms based on toy models using the operator product expansion as a starting point, and write the local fields in a position representation, since Fourier transforming to momentum space requires a well-behaved global structure of the space. A generalization of the Epstein-Glaser point-splitting procedure is used as the renormalization scheme. A consequence of this work is that there is in general no unique global vacuum state for the theory. The Wightman axiom requiring a unique invariant vacuum state is therefore completely rejected, while other axioms are generalized.<sup>3</sup> Since the reason for rejecting the vacuum axiom has to do with restriction to local regions, this lesson is compatible with local unitary equivalence and IR regularization: our idealization of infinite spacetime volume is again leading us astray. The hint here is that we will need to address this idealization in order to describe particle physics systems with spacetime curvature, and the suggestion by Hollands and Wald is that the formalism must be heavily modified to excise the vacuum concept altogether. As Belot (1998) has influentially argued, theory interpretation should be guided by the way that a given theory can be made compatible with other physical theories. Since the idealization of globally flat spacetime is believed to be false of our world, insights from attempts to fit QFT onto curved spacetime backgrounds are particularly illuminating. There are other reasons for thinking that the vacuum sector of LQFT is poorly understood—notably the cosmological constant problem (Koberinski 2021; Koberinski and Smeenk 2023)—so this solution strategy is interesting for reasons beyond just Haag’s theorem. For Haag’s theorem, the lesson is that we must restrict to local regions.

In all cases summarized here, we can understand the impossibility of the interaction picture as arising from overextending idealizations that are innocent for finite systems, but break down when one introduces continuum fields. The mathematics of infinite quantum systems requires more subtlety than finite systems. Despite other issues with AQFT, we see that Haag’s theorem can be avoided by modelling systems as having only finite extent, while retaining a continuum field description. Alternatively, one can weaken the notion of equivalence between interacting and free fields. Given that unitary equivalence is a more useful notion of equivalence elsewhere in quantum physics, while infinite spacetime extent is clearly an unphysical modelling idealization, the best conclusion is that Haag’s theorem is a pathology that emerges from requiring global equivalences between theories that are only meant to describe local regions. We should cautiously note, however, that the conclusion that infinite idealizations are *uniquely* responsible for Haag’s theorem *in all possible models* is too strong; it might turn out to be the case that counterexamples can be constructed, where a version of Haag’s theorem holds in finite spacetime regions. What is important, however, is that in the cases of physical interest Haag’s theorem is tied to these issues.

---

<sup>3</sup>See Koberinski (2021, Sec. 3) or Hollands and Wald (2010) for a more detailed summary of the Hollands and Wald approach.

## 4 Conclusions

What good is a no-go theorem for progress in the foundations of physics? For the project of interpreting our best theories and their implications for real-world systems, a no-go theorem provides a clear demonstration that some set of *prima facie* reasonable assumptions jointly lead to an unreasonable conclusion. Figuring out the various possibilities for modifying or rejecting assumptions, or living with the conclusion can provide deep insight into the physical implications of our best theories, and the idealizations and approximations necessary to link them to experiment. No-go theorems are often best formulated and analyzed in frameworks amenable to clear logical or mathematical presentations, where the necessity of each assumption is explicit. What we often learn from no-go theorems are better modelling techniques, or theoretical creativity for dealing with a supposedly unreasonable conclusion. This theoretical creativity can also lead to progress in constructing new theories.

For Haag’s theorem in particular, it seems to imply that a central pillar of perturbative calculation in LQFT—the interaction picture—is mathematically inconsistent. Careful attention to actual physics practice shows that this conclusion is too quick; regularization techniques otherwise required to obtain finite results in LQFT seem to do double duty in restoring the validity of the interaction picture. But the quantities of interest for particle physics are obtained after removing the regulators, so the physical significance of such a move is opaque. By moving to the AQFT formalism, we gain important insight into the nature of modelling idealizations for particle physics. In particular, we see that Haag’s theorem is a result of an infinite spacetime volume idealization, and disappears when one uses different formal tools to avoid this idealization. This sharpens the suggestion from LQFT, where it was unclear whether the UV or IR regulators were to blame, or if it was essential to move to a finite number of degrees of freedom. Additional insight into formulating AQFT on curved spacetimes suggests that the formalism will have to be fully local, and built from different quantities than vacuum expectation values. To answer the first question posed in the introduction, the interaction picture works because the apparent inconsistency is not actual, but due instead to taking an idealization too seriously.

Haag’s theorem typifies a good no-go theorem in physics. Once one makes the additional move to see AQFT and LQFT as complementary modelling frameworks, Haag’s theorem provides interesting physical insight into particle physics systems, and the space of theoretical possibilities that extend beyond this domain. When AQFT and LQFT provide convergent physical insight, we have better reason to believe that we have learned something about the actual physical systems under study. The axiomatic framework supplements insight gained from foundational analysis of LQFT to tell us that Haag’s theorem is not, strictly speaking, an artifact of a continuum field description, but is an artifact of an infinite volume idealization for relativistic systems. At a given scale, and to a given degree of approximation, a field description must be compatible with the existence of the interaction picture. We gain insight into how by combining interpretations from LQFT and AQFT.

### Conflict of interest statement

The author did not receive support from any organization for the submitted work, and has no competing interests to disclose.

## References

- Aizenman, Michael and Hugo Duminil-Copin (2021). “Marginal triviality of the scaling limits of critical 4D Ising and  $\lambda\phi_4^4$  models”. In: *Annals of Mathematics* 194.1, pp. 163–235. DOI: 10.4007/annals.2021.194.1.3. URL: <https://doi.org/10.4007/annals.2021.194.1.3>.
- Batterman, Robert W (2021). *A middle way: A non-fundamental approach to many-body physics*. Oxford University Press.
- Belot, Gordon (1998). “Understanding electromagnetism”. In: *The British Journal for the Philosophy of Science* 49.4, pp. 531–555.
- Duncan, Anthony (2012). *The conceptual framework of quantum field theory*. Oxford University Press.
- Earman, John and Doreen Fraser (2006). “Haag’s theorem and its implications for the foundations of quantum field theory”. In: *Erkenntnis* 64.3, pp. 305–344.
- Fraser, Doreen (2006). “Haag’s theorem and the interpretation of quantum field theories with interactions”. PhD thesis. University of Pittsburgh.
- Fraser, Doreen (2009). “Quantum field theory: Underdetermination, inconsistency, and idealization”. In: *Philosophy of Science* 76.4, pp. 536–567.
- Freeborn, David Peter Wallis, Marian Gilton, and Chris Mitsch (2022). *How Haag-tied is QFT, really?* URL: <http://philsci-archive.pitt.edu/21552/>.
- Haag, Rudolf (1955). “On quantum field theories”. In: *Kgl. Danske Videnskab. Selskab, Mat.-Fys. Medd.* 29.12, pp. 1–37.
- Haag, Rudolf (1958). “Quantum Field Theories with Composite Particles and Asymptotic Conditions”. In: *Physical Review* 112 (2), pp. 669–673. DOI: 10.1103/PhysRev.112.669.
- Hall, Dick and Arthur S. Wightman (1957). “A theorem on invariant analytic functions with applications to relativistic quantum field theory”. In: *Det Kongelige Danske Videnskabernes Selskab, Matematisk-fysiske Meddelelser* 31.5, pp. 1–41.
- Hollands, Stefan and Robert M Wald (2010). “Axiomatic quantum field theory in curved spacetime”. In: *Communications in Mathematical Physics* 293.1, pp. 85–125.
- Jost, Res (1961). “Properties of Wightman functions”. In: *Lectures on the Many-body Problems*. Elsevier, pp. 127–145.
- Klaczynski, Lutz (2016). “Haag’s theorem in renormalised quantum field theories”. In: *arXiv preprint arXiv:1602.00662*.
- Koberinski, Adam (2016). “Reconciling axiomatic quantum field theory with cutoff-dependent particle physics”. In: URL: <http://philsci-archive.pitt.edu/12496/>.
- Koberinski, Adam (2021). “Regularizing (away) vacuum energy”. In: *Foundations of Physics* 51.1, pp. 1–22.
- Koberinski, Adam and Doreen Fraser (2023). “Renormalization group methods and the epistemology of effective field theories”. In: *Studies in History and Philosophy of Science* 98, pp. 14–28.
- Koberinski, Adam and Chris Smeenk (2023). “ $\Lambda$  and the limits of effective field theory”. In: *Philosophy of Science* 90.2, pp. 454–474.
- Miller, Michael E. (2018). “Haag’s theorem, apparent inconsistency, and the empirical adequacy of quantum field theory”. In: *The British Journal for the Philosophy of Science*.
- Reed, Michael and Barry Simon (1975). “Methods of mathematical physics. Vol. II Fourier analysis, self-adjointness”. In.
- Reeh, Helmut and Siegfried Schlieder (1961). “Bemerkungen zur unitäräquivalenz von lorentzinvarianten feldern”. In: *Il Nuovo Cimento* 22.5, pp. 1051–1068.
- Ruelle, David (1962). “On the asymptotic condition in quantum field theory”. In: *Helvetica Physica Acta* 35, pp. 147–163.
- Streater, Raymond Frederick and Arthur S. Wightman (1964). *PCT, spin and statistics, and all that*. Princeton University Press.

- Swanson, Noel (2017). “A philosopher’s guide to the foundations of quantum field theory”. In: *Philosophy Compass* 12.5, e12414.
- Wallace, David (2011). “Taking particle physics seriously: A critique of the algebraic approach to quantum field theory”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 42.2, pp. 116–125.
- Wallace, David (2021). “The quantum theory of fields”. In: *The Routledge Companion to Philosophy of Physics*. Ed. by Eleanor Knox and Alistair Wilson. Routledge, pp. 275–295.