Still no peace on the lattice

Sébastien Rivat*

July 21, 2025

Abstract

The idea of using lattice methods to provide a mathematically well-defined formulation of realistic effective quantum field theories (QFTs) and clarify their physical content has gained traction in the last decades. In this paper, I argue that this strategy faces a two-sided obstacle: realistic lattice QFTs are (i) too different from their effective continuum counterparts even at low energies to serve as their foundational proxies and (ii) far from reproducing all of their empirical and explanatory successes to replace them altogether. I briefly conclude with some lessons for the foundations of QFT.

1 Introduction

Realistic quantum field theories (QFTs), which still figure as the most fundamental and empirically successful theories ever built, have been fraught since their inception with deep mathematical and conceptual issues. The foundational toll even seems to have grown out of control over time, moving from infinite energy levels and infinite probabilistic predictions to divergent perturbative series, divergent interaction parameters at a finite scale, inconsistent perturbative schemes, and ill-defined functional measures, to name only some of the toughest latecomers.

Physicists have designed two main strategies in response: (i) a pragmatic work-around strategy based on renormalization and effective field theory (EFT) methods; (ii) an axiomatic start-afresh strategy based on more advanced mathematical methods. The first has yielded approximate models whose empirical success remains unprecedented in the history of physics but whose mathematical and conceptual structure still lacks a proper foundation. For instance, the functional measure of realistic continuum EFTs is still ill-defined, strictly speaking. The second strategy has enabled physicists to

*Munich Center for Mathematical Philosophy, LMU Munich. Email: sebastien.rivat@lmu.de

identify candidate sets of axioms and well-defined mathematical structures for QFTs. But they are still yet to be shown to be instantiated by sufficiently realistic models. As it stands, although both strategies have contributed in their own way to the foundations of QFT, neither has succeeded in meeting all the legitimate demands of interested parties yet.[1]

There is nonetheless a promising middle path. The felt need to account more precisely for the mathematical structure of successful models, boosted by the widespread belief that QFT is probably not meant to deliver fundamental ones, has made increasingly attractive the idea of reformulating realistic QFTs as lattice models. As for renormalization and EFT methods, this strategy requires relinquishing some of their most cherished features: most obviously, their continuous space-time symmetries. But since these symmetries are unlikely to be fundamental and since there is not yet any ideal solution, there has been a growing recognition that lattice QFT currently provides a legitimate, if not the best, way to put realistic effective models on a perfectly firm mathematical and conceptual footing (e.g., Wallace, 2006; Duncan, 2012; J. D. Fraser, 2020b). Lattice QFT is indeed arguably the only game in town that has exact realistic interacting models and thus an adequate set of interpretative targets (cf. J. D. Fraser, 2020a).[2]

I argue in this paper that this middle-path strategy faces a two-sided obstacle. The most interesting side in my sense is theoretical: realistic lattice QFTs are just too different from realistic continuum EFTs even at low energies to serve as their foundational proxies, where, by 'foundational proxy,' I mean a mathematically well-defined reformulation of a theory (or model) that is sufficiently similar to its original formulation in the relevant respects to clarify its physical content (I provide more detail in section 4). We may of course bypass this obstacle by declaring that QFT was not at all the way we thought it was and endorsing lattice QFT as a new foundation for realistic effective QFTs, independently of its relationship to existing continuum EFTs. But we face again a serious albeit somewhat less interesting obstacle in this case: namely, lattice QFTs are far from reproducing all the empirical and explanatory successes of their effective continuum counterparts to replace them altogether. In fact, the very idea of using lattice QFT for this purpose would even sound incongruous to many lattice practitioners. Lattice QFTs have been designed to probe non-perturbative aspects of their continuum

[1]In fact, mixed methods involving mathematically more rigorous versions of renormalization and EFT methods have become increasingly popular among axiomatic and algebraic quantum field theorists, witness the development of the constructive field theory and causal perturbation theory programs (see, e.g., J. D. Fraser and Rejzner, 2024; Blum, forthcoming, for more historical detail). As we will see below, lattice QFT can be seen as a particularly successful instance of this trend.

[2]I use 'effective QFT' to refer to any kind of non-fundamental QFT model and 'continuum EFTs' for the subset of such models built on a continuum background.

counterparts, whether fundamental or effective (e.g., Creutz, 1985, p. 5; Degrand and Detar, 2006, p. 1). But they are rather impractical for the set of standard perturbative applications that underwrite the remarkable success of realistic QFTs (see, e.g., Capitani, 2003, for an introduction to lattice perturbation theory).

The paper is organized as follows. Section 2 briefly rehearses key aspects of continuum QFTs together with the set of traditional issues that affect them. Section 3 introduces lattice QFT, emphasizing, in particular, how it fully dissolves those issues. Section 4 provides a detailed account of how lattice QFTs differ from their effective continuum counterparts. Section 5 responds to various attempts to save the foundational role of lattice QFT. Section 6 concludes.

Three disclaimers are in order. First, the argument is largely independent of interpretative issues surrounding cut-offs and the existence of disanalogies between particle and condensed matter physics (e.g., Wallace, 2006, 2011; D. Fraser, 2009, 2011, 2020; Rosaler and Harlander, 2019). In particular, taking lattice QFTs seriously for foundational purposes does not require interpreting their lattice spacing overly strictly, say, as referring to some fundamental discreteness. As for continuum EFTs, we may perfectly well interpret the lattice spacing of a model as some unknown scale beyond which the effects of new physics become too significant for the model to be able to accommodate them.

Second, the argument is largely independent of the epistemic standing of the EFT framework. There are indeed very good theoretical and empirical reasons to believe that the most fundamental and empirically successful QFTs built so far are best formulated as effective theories, independently of their ill-defined mathematical and conceptual structure (see, e.g., Langacker, 2017, sec. 10.1; Isidori et al., 2024, pp. 2-5, for key limitations in the case of the standard model (SM) of particle physics). I only take issue with the additional step of taking our best continuum EFTs to be best formulated as lattice QFTs here.

Third, continuum EFTs form a diverse group including smooth momentum cut-off models, sharp momentum cut-off models, and dimensionally regularized models restricted to a UV-bounded segment of a renormalization group (RG) flow, to mention only a few (see, e.g., Bain, 2013; Rivat, 2025, sec. 2, for an introduction to continuum EFTs). I will take the latter as my reference point when comparing lattice QFTs with continuum EFTs in section 4. Diverse types of continuum EFTs may well exhibit different low-energy features too. But although I suspect that those differences (if any) are much less significant than for lattice QFTs, I will not have the space to argue for that here.

2 The hotbed of continuum QFTs

The most successful formulations of realistic QFTs, whether putatively fundamental or effective, rely on the path (or functional) integral formalism. To a large extent, this is also true of generalizations of standard QFTs, say, to non-equilibrium situations and gravitational contexts, as well as of the discretized versions below. I will accordingly start by presenting some of its key features, as a manner of introduction to the hotbed of issues that plague realistic QFTs.

Particle physicists are typically interested in computing the probability amplitude $\langle f|i\rangle$ for a given set of fields, say, a scalar field $\phi(x)$, to evolve from some initial configuration state $|i\rangle$ to some final configuration state $|f\rangle$. To obtain this kind of experimental quantity, they usually first compute the probability amplitude $\langle\Omega|\phi(x_1)...\phi(x_n)|\Omega\rangle$ for a more generic process where the field of interest, which is initially in its vacuum state $|\Omega\rangle$, transits into various configuration states across space-time $\phi(x_1)...\phi(x_n)|\Omega\rangle$ until it decays again into its vacuum state $|\Omega\rangle$. These generic probability amplitudes are referred to as "correlation functions" in the general case. As it happens, they can be derived from an even more elementary mathematical building block $Z[J]$ called the "generating functional," which encodes all the required dynamical information about the system and reduces to the so-called "path integral" once we ignore external sources J:

$$Z[0] = \int d[\phi(x)]e^{iS[\phi(x)]}, \tag{1}$$

where the measure $d[\phi(x)]$ specifies the distance between two arbitrarily close configurations and the action S encodes the dynamics of the system. Roughly put, the path integral represents the time-evolution of the system by attributing different weights e^{iS} to the different intermediary configurations that $\phi(x)$ may evolve into in between some initial and final states.

Despite the remarkable success of this formalism, realistic QFTs still remain affected by serious mathematical and conceptual issues. Perhaps the most ancient and well-known one, the so-called "problem of ultraviolet divergences," stems from the fact that most of the correlation functions obtained by applying perturbative methods to the path integral are infinite when the model of interest ranges over arbitrarily high energies. The renormalization program brings some relief (see, e.g., Butterfield and Bouatta, 2015; Rivat, 2019, for a philosophical discussion). But: (i) many perturbatively renormalized realistic QFTs still have at least some of their interaction parameters diverge either at some low- or high-energy scale (the so-called infrared and ultraviolet "Landau poles"); (ii) renormalized experimental quantities typically take the form of asymptotic perturbative series, which means (very roughly) that they diverge in a relatively well-behaved way. To this day, none of these more advanced issues has been fully solved by constructing more robust realistic models or using improved approximation methods. For instance, general considerations suggest that most of the perturbative series obtained from realistic QFTs are not Borel summable, which means that they cannot be uniquely matched to some exact function with the help of this rather natural resummation method (see, e.g., Fischer, 1997; Beneke, 1999, for more detail; Miller, 2021, 2023, for a discussion of the philosophical significance of asymptotic perturbative series).

But that is not all. Realistic QFTs are also strewn with a host of more abstract mathematical and conceptual issues that cast doubt on the well-standing of the QFT framework itself and its various approximation schemes (rather than on a specific set of realistic models). For instance, more abstract investigations into the structure of QFT have shown that the standard implementation of perturbative methods in the presence of interactions is inconsistent, strictly speaking. This result is enshrined in what came to be known as Haag's theorem (see, e.g., Earman and D. Fraser, 2006; Miller, 2018; Mitsch et al., 2024, for a philosophical discussion). The use of infinite-dimensional Hilbert spaces is also usually taken to raise important conceptual and interpretative issues (see, e.g., Ruetsche, 2011; Baker, 2016; Earman, 2020, for a philosophical discussion of unitarily inequivalent representations and non-separable Hilbert spaces).

The path integral formulation of realistic models also comes with its own mathematical and conceptual twists. Two are worthy of notice for models involving interacting bosonic fields. First, the path integral measure $d[\phi(x)]$ does not have a mathematically precise definition. The reason runs deep: little is known about the structure of the space of functions over which we integrate.[3] Second, the path integral Z contains a complex undamped exponential term $\exp(iS[\phi(x)])$ that oscillates widely for configurations $\phi(x)$ that are far from minimizing S. There is thus in general no reason to expect Z to converge.

As it happens, EFT methods have brought a welcome pragmatic solution to most of these issues. For instance, if we introduce a sharp momentum cut-off at low and high energies, we immediately obtain finite perturbative predictions. UV (resp. IR) Landau poles can also be easily avoided by imposing a sufficiently low-energy (resp. high-energy) boundary on renormalization-scale-dependent interaction parameters. And even if we do not restrict the range of realistic models so abruptly, or even explicitly, treating them as renormalized continuum EFTs already goes a long way toward taming their most threatening infinities. In particular, the traditional Landau pole sin-

[3] Strictly speaking, the same issue arises for fermionic fields. But it is less pressing since all the functional integrals involving Grassmann variables are finite.

gularities affecting realistic models seem to disappear once we introduce higher-order interaction terms (see, e.g., Djukanovic et al., 2018, for a discussion).

The EFT framework also cuts off the natural expectation of making exact predictions. Continuum EFTs are indeed designed to make more or less precise predictions by taking into account the contributions of a larger or smaller set of higher-order interaction terms (see, e.g., Georgi, 1993, p. 214, for a simple explanation). To make arbitrarily precise perturbative predictions, we need in principle to include every possible interaction term (since they are eventually all required when we consider higher orders in perturbation theory) and thus fix the value of an infinite number of independent parameters. Since this is impossible in practice, the threat of asymptotic series loses much of its bite in the EFT framework: there is just no reason to take seriously the full series for predictive purposes.

For all its merits, the standard continuum formulation of EFTs is still far from providing a principled answer to the most deeply entrenched issues affecting realistic QFTs. Consider again the issue of asymptotic perturbative series, which take the schematic form $A = \sum_n n!\ g^n A_n$ in the simplest cases, with g some interaction parameter. Imposing a finite cut-off on each sub-amplitude A_n does not undercut what typically makes those series divergent, to wit, the factorial contribution $n!$ at each perturbative order. There is thus in general no reason to expect the perturbative series derived from continuum EFTs to have better convergence properties than those of their perturbatively renormalizable counterparts—if anything, the new contributions arising from higher-order interaction terms at each perturbative order make the situation even worse. One may of course try to massage the divergent behavior of perturbative series by using Borel resummation techniques for instance. But similar types of ambiguities seem to arise in standard cases of EFTs too (see, e.g., Luke et al., 1995, for a discussion).

Perhaps even more problematically, continuum EFTs are still formulated by means of a continuum path integral with an ill-defined measure. To take the conceptually simplest example, suppose that the continuum EFT of interest is defined by separating its field variables according to some separation scale Λ, i.e., $\phi(x) = \phi_{<\Lambda}(x) + \phi_{>\Lambda}(x)$, and integrating out its high-energy field configurations $\phi_{>\Lambda}(x)$. The resulting path integral is still formulated with a similar continuum measure:

$$Z[0] = \int d[\phi_{<\Lambda}(x)] e^{iS_{\text{eff}}[\phi_{<\Lambda}(x)]}, \tag{2}$$

$$\text{with} \quad e^{iS_{\text{eff}}[\phi_{<\Lambda}(x)]} = \int d[\phi_{>\Lambda}(x)] e^{iS[\phi_{<\Lambda}(x), \phi_{>\Lambda}(x)]}.$$

Worse still, the very definition of the effective action S_{eff} is on unstable foundations, insofar as it is defined by means of a functional integral average over a continuum of high-energy field configurations.[4]

To sum up, continuum EFTs do bring some relief. Yet they still inherit some of the deepest mathematical and conceptual issues that affect realistic continuum QFTs. For those who like to keep their feet on firm mathematical ground and work with concrete realistic models, the lattice formulation of QFTs offers a way out, as we are now about to see (see, e.g., Montvay and Münster, 1994; Smit, 2002; Moore, 2003; Maas, 2020, for conceptually insightful introductions).

3 Lattice QFTs to the rescue

The key idea of lattice QFT is to reduce continuum quantum fields to finite-dimensional systems by assigning field variables only to every site and link of a discrete lattice of finite extent. The most common choice is to pick a periodic hypercubic Euclidean lattice with spacing a and size na, with n^4 the total number of lattice points, i.e., $\mathcal{M} = \{x | x_\mu/a \in [0, n[\}$.[5] In the case of our simple scalar example, the system is thus specified by assigning a field variable $\phi(x)$ at each lattice point x. The Fourier transform of $\phi(x)$ in momentum space, i.e., $\phi(p)$, takes values in the first Brillouin zone $\mathrm{B_Z} =]-\pi/a, \pi/a]^4$. For a finite-dimensional lattice, both $\phi(x)$ and $\phi(p)$ take values on finitely many position and momentum sites, respectively. And all the usual continuous derivatives and integrals in position and momentum space are replaced by differences and sums.

From there, the path integral takes a familiar form:

$$Z[a] = \int \Pi_x d\phi(x) e^{-S[\phi(x),a]}, \tag{3}$$

where the continuum integration measure is replaced by a discrete product of Lebesgue measures at each lattice point x. Although a different formulation is usually used for concrete applications, the discretized action for ϕ^4-theory

[4] Note that imposing a sharp momentum cut-off, either on the perturbative integral expression of correlation functions in momentum space or at the level of the path integral measure, is not the same as discretizing the model of interest (see, e.g., Rivat, 2019, pp. 9-10, for a toy example).

[5] Other choices are of course possible (e.g., a discretized space with continuous time, a random lattice). But a hypercubic Euclidean lattice is often privileged for computational purposes. For convenience, I will sometimes use a lattice of infinite extent $\mathcal{M} = a\mathbb{Z}^4$ to simplify equations in what follows.

can also be expressed in a familiar form:

$$S[\phi, a] = \frac{1}{2} \sum_{x,\mu} a^4 (\partial^f_\mu \phi \partial^f_\mu \phi + m^2 \phi^2) + \frac{\lambda}{4!} \sum_x a^4 \phi^4(x), \tag{4}$$

with $\partial^f_\mu \phi(x) = (\phi(x+\mu) - \phi(x))/a$ the forward derivative, λ some interaction parameter, and m the mass of the field. The action involves both a sum over every possible direction μ within each elementary hypercube and every discrete point x in the hypercubic lattice. As in standard QFT, we can use the discrete path integral to define a generating functional $Z[a, J]$ and derive correlation functions $\langle \phi(x_1)...\phi(x_n) \rangle$ on the lattice, either with traditional perturbative methods in the interaction parameter, other perturbative methods like the strong coupling expansion, or even non-perturbatively through numerical simulations (depending on the model and one's aims).

Although ϕ^4-theory is used as a component part of realistic QFTs, the real power of the lattice comes in full display in the context of lattice gauge theories. The latticization involves a new set of concepts in this case, two of which are particularly important. The first is the link variable $U(x, x+\mu)$, which links the value of a field, say, a scalar field or a spinor field, at a point x to its value at a neighboring point in the μ direction. We can then define a gauge field $A_\mu(x) = A^a_\mu(x) T^a$ as the generator of this elementary transformation, i.e., $U(x, x+\mu) = \exp(a A_\mu)$, with T^a the generators of the relevant Lie algebra used for a given gauge model. Note that A_μ lives on a link, strictly speaking. The second is the plaquette variable along a closed path. The most elementary plaquette variable defined around the smallest square of the lattice is given by:

$$U_p = U(x, x+\mu) U(x+\mu, x+\mu+\nu) U^\dagger(x+\nu, x+\mu+\nu) U^\dagger(x, x+\nu). \tag{5}$$

We can again define the field strength $F_{\mu\nu}(x)$ as the generator of this elementary transformation $U_p = \exp(a^2 F_{\mu\nu})$. And we can of course define plaquette variables around more complicated loops (Fig. 1).

Coming back to dynamical matters, the simplest and most common gauge-invariant action for a pure lattice gauge theory is the so-called "Wilson action:"

$$S[U] = \sum_p \beta \Big[1 - \frac{1}{d_F} \mathrm{Re}(\mathrm{Tr}\, U_p) \Big]. \tag{6}$$

The index p runs over arbitrary elementary plaquettes. The trace runs over gauge indices (e.g., $\mathrm{Tr}(F_{\mu\nu} F^{\mu\nu}) = F^a_{\mu\nu} F^{a\mu\nu}$, with $a = 1, ..., 8$ for $SU(3)$). d_F is the dimension of the fundamental representation of the gauge group and $\beta = 2 d_F / g^2$, with g some interaction parameter chosen so that $S[U]$ reduces

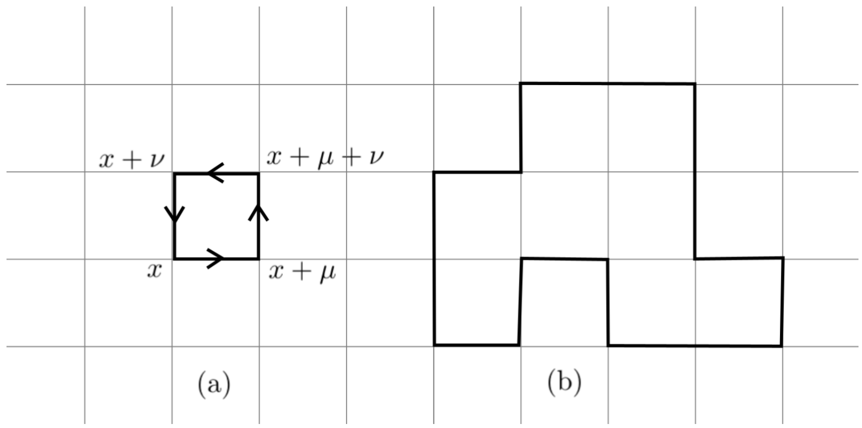

Figure 1: An elementary plaquette (a) and an arbitrary plaquette (b) in two dimensions.

to the standard pure gauge action in the (classical) continuum limit $a \to 0$:

$$S = \frac{1}{g^2} \sum_{x,\mu,\nu} a^4 \,\mathrm{Tr}(F_{\mu\nu} F^{\mu\nu}) + O(a^5). \tag{7}$$

Two remarks are in order at this point. First, the action is automatically invariant under local gauge transformations $U'(x, x+\mu) = \Lambda^{-1}(x) U(x, x+\mu)\Lambda(x+\mu)$ (since $\mathrm{Tr}\, U'_p = \mathrm{Tr}\, U_p$). Second, the link variables can be used to couple other field variables located at lattice sites:

$$Z[a] = \int \Pi_x d\phi(x)\ \Pi_{x,\mu} dU(x, x+\mu) e^{-S[U,\phi,a]}, \tag{8}$$

where $dU(x, x+\mu)$ is the Haar measure on the compact Lie group of gauge transformations. This path integral may be used to compute gauge-invariant quantities. Of particular interest is the Wilson loop, i.e., the expectation value of the trace of a plaquette.

Now, for the good news, lattice QFT allows us to dissolve every single mathematical and conceptual issue we have been worried about so far. As we will see in section 4, there is a cost to that. But I will bracket it for now.

Starting with the traditional issues, then, all the momentum integrals involved in (perturbative) correlation functions reduce to discrete sums with a high-energy cut-off π/a. (If needed, we may also introduce a low-energy regulator to eliminate infrared divergences in perturbation theory.) Every (perturbatively) renormalized interaction parameter is also either fixed at the lattice scale a or depends on some scale whose interval is appropriately bounded. And so every (perturbative) prediction (at a given order) is finite in lattice QFT.

The more fundamental issues plaguing realistic QFTs disappear too. (i) The premises of Haag's theorem are simply not satisfied on a lattice of finite extent. (ii) The traditional issues arising from the putative relevance of non-separable Hilbert spaces and infinite-dimensional unitarily inequivalent representations are also irrelevant in this setting. (iii) Lattice QFT opens the door to non-perturbative numerical computations and new kinds of perturbative methods that are better behaved than the standard ones. Of particular importance is the strong coupling expansion, which has a finite radius of convergence (see, e.g., Itzykson and Drouffe, 1989, sec. 6.3.1). (iv) The path integral measure in lattice QFT takes the form of a countable product of well-defined Lebesgue measures $d\phi(x)$ for scalar fields, Berezin measures $d\psi(x)$ for spinor fields, and Haar measures $dU(x, x+\mu)$ for link variables. Because we integrate over a compact Lie group instead of a non-compact Lie algebra in the latter case, we do not even need to partially fix the gauge to avoid divergences originating from gauge redundancies. (v) The standard Euclidean formulation (which is also sometimes used in the continuum QFT setting) ensures that the exponential term $\exp(-S[U, \phi])$ is dampened for field configurations (U, ϕ) that are far from minimizing S. All in all, lattice QFTs succeed with brio to escape all the trouble we encounter with their continuum counterparts.

Does that mean that we should take lattice QFT seriously for foundational and interpretative purposes? In the wake of the success of EFTs, philosophers have become increasingly inclined to recognize that lattice methods provide a legitimate way to understand the structure of realistic effective QFTs (e.g., J. D. Fraser, 2018, 2020b; Rosaler and Harlander, 2019; Williams, 2019).[6] As James Fraser puts it, assimilating cut-off QFTs with lattice QFTs here (cf. J. D. Fraser, 2020b, p. 282):

> The central problem here was the lack of a clear answer to the question of what empirically successful QFTs are—both mathematically and physically. We saw in Section 14.3, however, that it is possible to precisely define the path integral for the partition function and explicitly construct realistic QFTs as mathematical models if the degrees of freedom of the field associated with arbitrarily large energies and momenta are removed via the cutoff. This provides a non-perturbative characterization of QFT which has a crucial advantage over the axiomatic systems discussed in Section 14.2; we can explicitly write down cutoff formulations of empirically successful QFTs, and the standard model in particular. Furthermore, we have seen good reasons to regard these cutoff models as conceptually respectable, and empirically suc-

[6] Despite my preference for effective continuum fields, I also now think that I was too quick when appealing to scalar lattice models in (Rivat, 2021, secs. 4-5).

cessful, theories in their own right. (J. D. Fraser, 2020b, pp. 285-6)

To be fair, Fraser does not explicitly endorse lattice QFT as the correct foundational setup for understanding realistic effective QFTs. But he does seem to acknowledge its legitimacy. We may indeed perfectly well formulate realistic four-dimensional gauge field theories on a lattice and extract empirical numbers out of them without infringing on the highest standards of mathematical rigor and conceptual clarity.

Among all, David Wallace seems to come the closest to endorsing the lattice formulation of QFTs as the best current way of putting realistic effective QFTs on a perfectly secure mathematical and conceptual footing (see, esp., Wallace, 2006, p. 48). He is of course well aware of the diversity of existing regularization schemes and of the varying degree of mathematical well-definedness exemplified by different kinds of cut-off QFTs. He is also far from taking the claims that any particular cut-off model makes at its cut-off scale seriously. In fact, he does not even take the lattice formulation of a given QFT model (or the discretized version he outlines in 2006, p. 48) to be the correct way to formulate it. If anything, the correct definition of realistic effective QFTs will be provided by a mathematically well-defined superseding theory.[7] And in the absence of such a theory, the second-best choice is to define a realistic effective QFT through the equivalence class of cut-off QFTs with a similar low-energy structure that one may define for the model at stake.

But Wallace still seems to be ready to grant that some regularizations, including the standard lattice regularization and the discretization procedure he picks in (2006), provide a maximally robust mathematical characterization of realistic effective QFTs. This suggests, in turn, that the subset of lattice QFTs for a given model furnishes us with the mathematically cleanest representatives of the equivalence class of cut-off QFTs available for that model. And insofar as lattice QFTs enable us to put any such equivalence class on a perfectly firm mathematical and conceptual footing, they also seem to provide a privileged interpretative standpoint to understand the low-energy physical content of realistic effective models (until we find a mathematically well-defined and empirically successful superseding theory).[8]

[7]Note that this is a rather strong desideratum. We may well find realistic superseding theories that are as badly defined as our best current QFTs.

[8]One may wonder whether Wallace (2006, p. 48) is genuinely committed to any strict form of latticization or discretization, and may not easily fall back to more flexible forms of regularization, which remove all but a finite number of momentum modes $\phi(p)$—or even merely eliminate high- and low-energy ones—while leaving the continuum background structure of the regularized model intact. Although this is certainly worth exploring, I am inclined to be skeptical toward this strategy in its current shape. (i) The regularization

Now, as I will argue in sections 4-5, singling out lattice (or discretized) QFTs for such foundational and interpretative purposes is too quick. Even on this kind of sophisticated account, there is a principled theoretical obstacle to taking lattice QFTs to form a representative of continuum EFTs: namely, lattice QFTs are just too different from their effective continuum counterparts even at low energies to be put in the same low-energy equivalence class. We may of course take those differences seriously and concede that lattice QFT provides us with an altogether new foundation for realistic effective QFTs. But as we will see in section 5, there is a serious obstacle waiting for us in this case too.

4 A radical departure

I will now examine more closely how the latticization of a continuum QFT distorts its structure across scales.

4.1 Preliminary remarks

We first need to clarify the idea of using lattice QFT to put continuum EFTs on a robust mathematical and conceptual footing. There are several options here.

The first and perhaps the most intuitive one is to appeal to the traditional axiomatic account: to put a theory (or a model) on a robust mathematical and conceptual footing is to show that it instantiates a mathematically precise and consistent set of definitions and axioms. In practice, we could try to use lattice QFTs to identify the set of definitions and axioms instantiated by their effective continuum counterparts. This would give us, in turn, a precise and general characterization of the class of systems they depict. Yet it is fairly clear that continuum EFTs and lattice QFTs do not have the same axiomatic structure, witness their different background space (I will provide more detail below). We could try to show instead that they instantiate equivalent sets of definitions and axioms. Yet, to establish any sort of equivalence theorem, we would first need to find the set of definitions and axioms instantiated by continuum EFTs, which, if successful, would directly undermine the foundational relevance of lattice QFT. Perhaps lattice models

would need to be concretely implemented for realistic models to properly assess its merits, both theoretical and practical. (ii) There are preliminary reasons to believe that the regularized models would be worse than their lattice counterparts. If we keep a continuum of field configurations, either in position or momentum space, we still have an ill-defined continuum path integral measure. Inversely, if we keep only a finite number of momentum or position modes in some unruly fashion, the resulting models will probably have even fewer symmetries than their lattice counterparts and depart even more from continuum EFTs.

should then be used to identify the set of definitions and axioms that their effective continuum counterparts approximately instantiate. This strategy, however, seems to be overly elusive. After all, an inconsistent theory may well approximately instantiate a consistent set of axioms. We might for instance think of a QFT model whose predictions violate only slightly some unitarity bound (see, e.g., Schwartz, 2013, sec. 24.1.5, for an introductory account). Thus, on this refined account, even if we succeed in clarifying the axiomatic structure of lattice QFTs, we might still be far from having a clear view on the foundational status of their effective continuum counterparts.

Another option is to appeal to the traditional reductionist account: to put a theory (or a model) on a robust mathematical and conceptual footing is to show that it can be derived from a mathematically well-defined and consistent theory (or model) with the help of auxiliary assumptions and approximations. There is no need to excavate the axiomatic structure of lattice QFTs and continuum EFTs in this case. Rather, the goal is to gain some confidence in the foundational standing of continuum EFTs by rederiving them, perhaps only approximately, from their mathematically well-defined lattice counterparts, and clarify their physical content by extracting some partial and approximate interpretative targets out of those lattice models. Yet it seems that this strategy defeats once again its purpose. To recover the continuum EFT of interest, we would first need to find a well-defined continuum limit for the relevant lattice QFT and then restrict again its domain, say, by integrating out its high-energy continuum field configurations. The real foundational work would thus be done in this case by the putatively fundamental QFT obtained in the continuum limit, thereby relegating the role of lattice QFT to that of a mere heuristic ladder to be thrown away once the derivation is complete.

A better option in my sense is to take a lattice QFT to provide a mathematically well-defined reformulation of a continuum EFT (see Hunt, 2024, forthcoming, for recent work on the nature and value of reformulations in physics). This does not require us to know anything about their axiomatic structure or to be able to derive them from each other. Rather, the goal of a foundational reformulation is to replace problematic component parts of an input theory (or model) with mathematically well-defined ones and use the output theory (or model) to clarify its physical content. In practice, we should thus use lattice QFTs to show that continuum EFTs are free from contradiction and mathematical ambiguity once properly reformulated.

Now, of course, the reformulation may affect more or less the mathematical structure of the input theory. As we have seen, the output and input theories presumably do not need to be theoretically equivalent or satisfy the same definitions and axioms. In particular, insofar as lattice QFTs and continuum EFTs are not meant to be fundamental, they may well depart radically from each other at the scale of a lattice spacing or a momentum cut-off. Yet, if the goal is ultimately to use lattice QFTs to clarify the physical content of their effective continuum counterparts and thus genuinely treat them as "foundational proxies," we should presumably require the input and output theories to be sufficiently similar to each other in the relevant respects (I briefly discuss alternative strategies at the end of section 5).

Finding the right notion of similarity for theories (or models) is nonetheless far from trivial and usually recognized to be context-dependent (e.g., Rosaler, 2015; Fletcher, forthcoming). There are two aspects to disentangle in our case: (i) clarify the sense in which the mathematical structures of the two models to be compared are similar to each other; (ii) identify the set of mathematical structures (and properties) that are physically significant and thus underwrite the extent to which the two models have approximately the same physical content.

Regarding (i), the similarities between the mathematical structures of lattice QFTs and continuum EFTs are well captured by two kinds of relations.

First, the models involve mathematical structures like complex functions for which the standard Lebesgue measure is rather natural. We may thus safely speak in this case of numerical approximation in the relevant range or domain, and talk of approximate truth for the corresponding physical statements. As we will see below, the ideal scenario in this case arises when a mathematical structure of a lattice QFT takes the form of a perturbative approximation in the lattice spacing parameter a of a mathematical structure of a continuum EFT, with an infinite amount of structure whose contributions become negligible for physical situations characterized by a sufficiently low energy scale $k \ll 1/a$.

Second, the models also contain more abstract mathematical structures for which it is more sensible to use a broader notion of approximate isomorphism. We may for instance wish to compare a discrete lattice with a continuum manifold, a discrete symmetry with a continuous one, or a Riemann sum with a continuous integral. These mathematical structures may well give rise to good numerical approximations, as when the value of a large Riemann sum approximates well the value of a continuous integral. But the underlying structures behind these quantities need not be approximately isomorphic to each other, as for the number of elements in a Riemann sum with respect to the continuum of values over which we integrate. The notion of approximate isomorphism may of course be further specified depending on the mathematical structure of interest. But in general, for the notion to provide an adequate measure of similarity, the approximate isomorphism arguably needs to range over a sufficiently comprehensive set of elements and preserve a sufficiently important number of relations and properties between these elements (see, e.g., Rivat, 2021, sec. 5, for a discussion).

Regarding (ii), we can again identify two sets of physically significant mathematical structures and properties. First, the path integral formalism makes it rather natural to interpret the basic physical content of a model in terms of its background space, degrees of freedom, dynamics, and correlation functions. Second, this basic physical content typically possesses physically significant higher-order structures and properties, whether we speak of symmetries, locality, analyticity or hermiticity for instance. In both cases, there is of course some leeway about how to best carve out the physical content of lattice QFTs and continuum EFTs. In particular, a more austere interpretation in terms of the numerical information encoded in correlation functions is a live option (see, e.g., J. D. Fraser, 2018; Ruetsche, 2020; Rivat, 2021, for a discussion). However, as I will briefly argue in section 5, this option seems to make continuum EFTs interpretatively irrelevant and thus to bring us back to the idea of taking lattice QFTs to replace them altogether.

On this rather natural account, then, we may portray the success case as follows: although lattice QFTs and continuum EFTs may differ radically from each other at the scale of a lattice spacing or momentum cut-off, most of their physically significant mathematical structures become approximately equal or isomorphic to each other and are governed for the most part by the same physically significant principles and properties at low energies, without any ad hoc intervention on our part (on pain of comparing the continuum EFT of interest with a different model).

To give a straightforward example of a successful low-energy approximation, consider the continuum effective model of a massless scalar field $\phi(x)$ with a Gaussian cut-off. The kinetic term takes the form $\partial_\mu\phi(x)\exp(-\partial^2/\Lambda^2)\partial^\mu\phi(x)$. High-energy perturbative contributions are accordingly damped by the modified propagator $\exp(-p^2/\Lambda^2)/p^2$. Now suppose that we define a continuum effective model with a less "peaked" Gaussian cut-off by using the kinetic damping factor $\exp(-\partial^4/\Lambda^4)$. The two models have exactly the same background space, degrees of freedom, and higher-order properties. The only difference lies in higher-order contributions to the dynamics and correlation functions. Once expanded, the more and less "peaked" kinetic damping factors $\exp(-\partial^2/\Lambda^2)$ and $\exp(-\partial^4/\Lambda^4)$ indeed give rise to an infinite series of different contributions to higher-order dynamical terms $\partial_\mu\phi(x)(\partial^2/\Lambda^2)^n\partial^\mu\phi(x)$ ($n \geq 1$) (which are in principle already included in the effective dynamics). Since all these different contributions become irrelevant at low energies, we seem to be justified in concluding that the two effective models have approximately the same physical content in this regime.

Now, the situation certainly becomes more thorny when comparing other kinds of continuum EFTs, say, a sharp cut-off and a dimensionally regularized model. Gauge and translation invariance are particularly worrying for sharp cut-off models. But even in such cases, the continuum EFTs of interest seem to be much more similar to each other at low energies than they are to lattice or discretized QFTs. As already emphasized, I unfortunately do not have the space to defend this claim here. Since lattice QFT has been deemed to enjoy a special foundational role, I will restrict myself to the claim that lattice models are too dissimilar from their effective continuum counterparts at low energies to serve as their foundational proxies (using the case of dimensionally regularized models as my reference point). As we will see, we can certainly obtain a number of important low-energy quantities that are in close numerical agreement with each other. Yet the latticization also leads us to chop off an infinite amount of physically significant mathematical structure and violate key physical principles down to arbitrarily low energies.

4.2 Basic physical content

Let us first focus on the basic physical content of lattice QFTs and continuum EFTs.

Starting with background spaces, the four-dimensional Minkowski space-time of a continuum QFT is replaced by a flat hypercubic lattice space with periodic boundary conditions in standard applications. Needless to say, the Euclideanization, discretization and compactification procedures make the two spaces nothing alike for any non-zero value of a, whether we speak of their cardinality, topology or metric (including the space-time split and related causal structures). The spatio-temporal structures of a lattice system and its effective continuum counterparts thus remain radically different from each other, no matter how we coarse-grain both types of systems. We may of course discretize Minkowski space-time in different ways, some of which may look at first sight more physically significant than others, as when we keep time continuous in the Hamiltonian formulation of lattice QFTs (Kogut and Susskind, 1975). But they tend to be more cumbersome and not as successful as the standard Euclidean discretization.[9]

The two types of models are also radically different with respect to their physical degrees of freedom. Most obviously, lattice QFTs involve only a finite number of degrees of freedom compared to their effective continuum counterparts. The space of lattice field configurations is likewise nowhere approximately isomorphic to the space of low-energy field configurations of a continuum EFT for any non-zero value of a (or ak). And this is true even if we impose both a low-energy and high-energy sharp momentum cut-off in the path integral measure of a continuum EFT: the resulting model still involves a continuum of variables on the space-time manifold as well as a

[9]One might think that those spatio-temporal differences are ultimately physically insignificant since there is little reason to take seriously the flat background space-time structure of relativistic QFTs in light of quantum gravity. As we will see in sections 4.3 and 5, however, even such seemingly innocuous differences do have a serious impact on a variety of other physically significant mathematical structures.

continuous space of field configurations (if we put issues pertaining to its mathematical definition aside).

Yet the difference is not just quantitative. Discretizing a model also has a significant impact on the selection of physically relevant dynamical variables. As we saw above, the most elementary gauge invariant structures of pure lattice gauge models are plaquette variables U_p (once we take the trace). Likewise, more general models make irreducible use of intrinsically non-local variables $U(x, x+\mu)$ that live on intermediary lattice links and whose values lie on the complex circle in contrast to the real-valued local gauge field variables of their effective continuum counterparts.

Those basic differences have important repercussions for the states of the systems studied too. I will speak in more detail about irreducible particle-state representations below. For now, the brief presentation of lattice QFTs outlined in section 3 is sufficient to make the point. For pure lattice gauge models, the only physical states are the so-called "glueball" states obtained by applying gauge-invariant operators to the vacuum state $|0\rangle$ (e.g., $\mathrm{Tr}(U_p)|0\rangle$ for an elementary plaquette state). For more general models, the physically salient configurations take the form of open and closed tetris-like configurations built out of site and link variables, as illustrated in Fig. 1, which makes lattice QFTs more akin to string theories than local field theories (see, e.g., Kogut and Susskind, 1975, sec. V; Montvay and Münster, 1994, sec. 3.6.1, for more detail).

Regarding the dynamics and correlation functions of lattice QFTs, there is a well-deserved respite: they ultimately do furnish a good numerical approximation to their effective continuum counterparts at sufficiently low energies. But as we will see below and in sections 4.3-5, this comes at the price of ad hoc fine-tuning maneuvers and significant higher-order physical differences.

Consider first the dynamics of lattice QFTs. A first pass at a lattice action suggests that it is very different from its effective continuum counterparts. (i) The Minkowski action is replaced by a Euclidean one in standard applications, whether it is directly posited or obtained by analytic continuation from its Minkowski counterpart ($t \to -it$). (ii) The continuum Euclidean action over local field variables is replaced by a discrete sum over local site variables and non-local link variables.[10] (iii) The lattice action typically involves non-polynomial interaction terms like $\phi^\dagger(x)\exp(iaA_\mu(x))\phi(x+\mu)$ instead of polynomial interaction terms like $\phi^\dagger(x)A_\mu(x)\partial^\mu\phi(x)$. (iv) The lattice action contains an infinite number of additional interaction terms arising from the use of differences instead of derivatives and from the weakly constrained symmetric structure of interaction terms on the lattice. In par-

[10] Note that we are also chopping "most" of the space of monomials of fields and their derivatives across space-time by taking into account only those defined on lattice sites and links.

ticular, the homogeneous Lorentz group $SO(3,1)$ forces us to keep only interaction terms with indices summed in pairs. By contrast, the hypercubic group $H(4)$ allows us to use interaction terms with indices summed in even numbers like $\sum_\mu (\partial_\mu \phi)^4$ and $\sum_\mu \phi(\partial_\mu^4 \phi)$ for instance (see, e.g., Moore, 2003, p. 8, for more detail). (v) The lattice action allows for a greater amount of operator mixing, including between operators of different dimensions (see, e.g., Capitani, 2003, sec. 14; Degrand and Detar, 2006, sec. 16.2.4). Depending on the model, we may thus be forced to fine-tune the parameters of relevant interaction terms generated upon renormalization in order to match a lattice action to its effective continuum counterparts.

Upon closer examination, however, we still seem to be able to erase at least the dynamical differences (iii)-(v) at sufficiently low energies with enough fine-tuning. We can indeed first re-express the lattice action in terms of site and gauge variables and make it closely resemble its effective continuum counterparts (in Euclidean space) with additional dynamical lattice artifacts organized in terms of increasing powers of the lattice spacing a. Then, it is easy to show that the dynamical lattice artifacts originating from the lack of spatio-temporal constraints all become negligible at arbitrarily low energies $ak \to 0$ for sufficiently well-behaved configurations. This is in part due to the fact that differentiable functions and their derivatives are well approximated by differences of functions for a small lattice spacing. As we will see in sections 4.3 and 5, the situation is more complicated with respect to other dynamical artifacts, especially for those that are relevant at low energies and require fine-tuning. But even for models involving such artifacts, there is still a clear sense in which the lattice action is similar to its effective continuum counterparts in the relevant regime: the sets of dynamical terms in the two actions become increasingly approximately isomorphic to each other as we neglect and fine-tune increasingly many dynamical lattice artifacts for arbitrarily small ak. Schematically:

$$S_{\text{latt}} = \underbrace{S_{\text{eff}}}_{\substack{\text{partly}\\ \text{fine-tuned}}} + \underbrace{S_0' + \frac{1}{a}S_1' + ... + \frac{1}{a^4}S_4'}_{\text{fine-tuned}} + \underbrace{aS_1 + a^2 S_2 + ...}_{\text{negligible}}, \tag{9}$$

where the parameters of S_{eff} include both relevant and irrelevant a-dependent renormalization contributions and S_i, S_i' include new dynamical lattice artifacts not present among S_{eff}'s interaction terms.[11]

[11]For instance, in non-abelian gauge theories, the action includes relevant dynamical lattice artifacts like a gluon mass term $g^2(A_\mu^a)^2/a^2$, which arises from the gauge-invariant measure (e.g., Capitani, 2003, sec. 5.2.1; Maas, 2020, pp. 90-1). We also have relevant perturbative contributions arising from irrelevant dynamical artifacts, say, an interaction term involving two quark fields and two gluon fields whose $O(a)$-dependence at tree level is compensated by $O(1/a)$-contributions from loop integrals. More generally, all the relevant and irrelevant dynamical lattice artifacts are required to maintain gauge

At the risk of belaboring the point, the situation is similar for correlation functions. At first sight, the latticization introduces again important structural differences. (i) The correlation functions are defined over a discrete Euclidean space for which no time ordering is required (or drastically modified, as when states are defined on successive Euclidean slices). (ii) The correlation functions provide only information about the average values of field products over a finite set of possible points. (iii) The Feynman rules become rather unusual if we use lattice QFTs in the perturbative setting: propagators and vertices inherit a complicated pattern of trigonometric dependence, and new vertices associated with dynamical lattice artifacts enter the scene too (see, e.g., Montvay and Münster, 1994, secs. 2.2.2, 3.3.1, 5.1.5, for some examples).

Yet, as for the lattice action, there are good reasons to believe that lattice correlation functions ultimately have approximately the same values as their effective continuum counterparts in the perturbative regime, at least at sufficiently low energies and modulo some heavy fine-tuning. Once relevant contributions are absorbed by adjusting the value of bare parameters, all the remaining perturbative contributions indeed become negligible or converge towards their effective continuum counterparts at low energies. To take the simplest example, the lattice propagator of a massive scalar field, which corresponds to a two-point non-interacting correlation function, converges to the continuum propagator at low energies $ak \rightarrow 0$:

$$\frac{1}{4\sum_{\mu}\frac{1}{a^2}\sin^2(\frac{ak_\mu}{2})+m^2} \rightarrow \frac{1}{k^2+m^2}, \tag{10}$$

with $k_\mu \in]-\pi/a, \pi/a]$. In the non-perturbative regime, by contrast, the situation is less clear since we do not yet have good non-perturbative control over realistic continuum EFTs. It is thus harder to assess for instance whether the correlation functions of the effective version of continuum quantum chromodynamics naturally match with their lattice counterpart at low energies.

To sum up, the basic low-energy physical content of lattice QFTs and continuum EFTs is radically different from each other. We can still use the dynamics and correlation functions of lattice QFTs to approximate that of their effective continuum counterparts at low energies. But there are good reasons to think that the similarity at play in this case is largely numerical. Not only is the mapping highly partial and subject to ad hoc fine-tuning maneuvers. But as we are now going to see, there are also many significant higher-order properties of the dynamics and correlation functions that are lost at arbitrarily low energies.

invariance in perturbation theory.

4.3 Physical principles

On the face of it, lattice QFTs violate almost all the sacred principles of physics, whether we speak of locality, space-time symmetry principles, or energy-momentum conservation laws. But we need to discuss whether this is a real problem at low energies.

Let me first say a brief word about locality, causality, and analyticity. (i) As we saw above, lattice QFTs make essential use of non-local link variables and non-local interaction terms, even in what is usually deemed the non-interacting part of their effective continuum counterparts (e.g., the term $\partial_\mu^f \phi \partial_\mu^f \phi$ in Equation 4). (ii) Lattice QFTs violate standard relativistic causal principles prohibiting spacelike physical processes and superluminal propagation for any non-zero value of a. In particular, neighboring scalar and spinor degrees of freedom seem to affect each other at a distance. The traditional microcausality condition of continuum QFTs does not make any sense either, strictly speaking. The change of background structure indeed requires a new metric δ_{ij} ($0 \leq i, j \leq 3$), which undermines any kind of meaningful distinction between spacelike and timelike separation (besides affecting the causal structure of events). (iii) Insofar as lattice models involve discrete functions, the discretization also makes us lose analyticity. Yet, in spite of it all, there is still an important sense in which the lattice action remains "local:" namely, the kinetic operator in position space decays sufficiently fast with distance, in the sense that it is bounded by some function $C \exp(-\gamma|x|)$, with C and γ some constants.[12]

Next, the latticization of a continuum QFT drastically affects its space-time symmetries, no matter how small a (or ak) is. Continuous translations (R^4) become discrete (Z_n^4): we can translate the system by one unit of the lattice spacing at a time with periodic boundary conditions. Time reversal invariance is replaced by reflection positivity, which must hold for both site and link reflections. The switch from Minkowski space-time to the continuum Euclidean space forces us to replace the homogeneous Lorentz group $SO(3,1)$ by the orthogonal group $SO(4)$. The latticization, in turn, forces us to replace $SO(4)$ by the hypercubic group $H(4)$, which consists of discrete block rotations by $\pi/2$ and reflections. To be sure, the remnant discrete group of lattice transformations constitutes a subgroup of its continuous Euclidean counterpart $R^4 \rtimes SO(4)$. There is also no issue with restoring $R^4 \rtimes SO(4)$ by taking $n \to \infty$ and $a \to 0$. But this does not alter the fact that the spatio-temporal symmetric structures of a lattice QFT and

[12] This brief warning is of course far from doing justice to the topic of locality, causality, and analyticity in non-relativistic quantum theories. Particularly worthy of notice here is the existence of generic bounds on the maximal speed at which information propagates (see, e.g., Nachtergaele et al., 2019, for a review). I am thankful to Benjamin Feintzeig for bringing this to my attention.

its effective continuum counterparts are far from being approximately isomorphic to each other for any non-zero value of a. Chopping "most" of the continuum structure of the background space indeed forces us to chop as well "most" of the continuum structure of its space-time symmetries no matter how coarse-grained the lattice system is.

This has two physically salient implications besides changing the transformation properties of the dynamics and correlation functions and making them less constrained (as we saw in section 4.2).

First, the loss of spatio-temporal symmetric structure leaves us with at best a partial energy-momentum conservation law, i.e., energy-momentum is conserved only up to $2\pi/a$ for a lattice with spacing a. Note that this does not mean that energy-momentum is approximately conserved at sufficiently low energies. Any such partial conservation law indeed allows for arbitrarily large violations of energy-momentum by $2n\pi/a$. Concretely, this means that, on a lattice, low-energy incoming particles with $k \sim 0$ can in principle give rise with high probability to an arbitrary even number of outgoing high-energy particles with $k \sim \pi/a$.

Second, the traditional classification of particles in terms of infinite-dimensional irreducible representations of the Poincaré group breaks down. We can of course distinguish between different kinds of lattice entities by appealing to the irreducible representations of the lattice symmetry group. But there are significant physical differences in this case too. For a start, the definition of momentum eigenstates on a Euclidean lattice of spacing a and finite extent na implies that: (i) the energy spectrum is discrete and bounded, with a smallest momentum $2\pi/na$ and a largest momentum π/a in any direction; (ii) the mass of a particle becomes a trigonometric function of its momentum. Worse still, since $H(4)$ is compact in contrast to $SO(3,1)$, we only have a finite number of irreducible representations on the lattice to account for known particles with non-zero spin or helicity. We may try to match this finite set to the infinite set available on the continuum. In practice, this is easier to do with a three-dimensional spatial lattice since the discrete cubic group of rotations $H(3)$ constitutes a subgroup of the group of spatial rotations in Minkowski space-time. But even in this simplified case, mapping the finite set of irreducible representations of $H(3)$ to those of $SO(3)$ still gives rise to serious interpretative ambiguities (see, e.g., Montvay and Münster, 1994, pp. 153-5; Maas, 2020, sec. 2.9.1.3, for a discussion). For instance, if we want to account for a massive continuum particle with a given spin, we are typically forced to decompose its reducible continuum representation into distinct irreducible lattice representations and thus "split" it into distinct lattice species. In practice, this means that we may for instance have to attribute different masses to lattice particles that have otherwise the same mass on the continuum (e.g., Montvay and Münster, 1994, pp. 154-5).

What about internal symmetries? The good news first: gauge symmetries are exactly preserved in lattice QFTs. One might even say that the lattice's *raison d'être* is to obtain a well-defined non-perturbative gauge-invariant formulation of QFTs. We do not need to fix partially the gauge in the path integral to obtain a quantized model. We do not need ghost fields and more sophisticated symmetries like BRST—although we can of course formulate perturbative versions of lattice gauge theories and recover all of this structure if needed. And the same assessment goes for other internal global symmetries like $O(N)$ for N-component scalar fields.[13]

But all is not perfect either. Physicists have found that there are principled reasons to believe that chiral invariance cannot be exactly implemented on the lattice without giving up on some other important principle. The issue is best introduced through the so-called "fermion doubling problem:" namely, the naïve formulation of lattice QFTs involving fermionic fields displays, as a matter of principle, redundant fermionic particle species (the so-called "fermion doublers" or "mirror fermions"). We can already see this generic pattern in the dispersion relation of a massless fermionic field in a two-dimensional lattice model (see, e.g., Tong, 2018, sec. 4.3.1, for a physically intuitive account):

$$E(k) = \frac{1}{a} \sin(ak). \tag{11}$$

This dispersion relation has two zeros in the first Brillouin zone, i.e., two distinct massless fermionic species with opposite chiralities for $k = 0$ and $k = \pi/a$. This result generalizes to any dimension d: every fermion has $2^d - 1$ partners, with the same number of left-handed and right-handed fermions in total. The underlying reason comes from the fact that the standard kinetic term for fermionic fields involves only first-order derivatives compared to scalar and gauge fields. This leads to the problematic dependence of the dispersion relation on $\sin(ak)$, as opposed to $\sin^2(ak/2)$ for scalar and gauge fields, which does not give rise to zeros at the edges of the Brillouin zone.[14]

[13] See, e.g., Dougherty (2021) and Rivat (forthcoming) for philosophical discussions related to perturbative continuum gauge theories. As it turns out, lattice gauge theories actually have more global gauge symmetries than their effective continuum counterparts. For instance, for the gauge group $SU(N)$, the lattice action is invariant under the center of the gauge group Z_N. Otherwise, more advanced symmetries like supersymmetry, which can be seen as "mixing" internal and external symmetries, are not preserved by standard latticization procedures (see, e.g., Bergner and Catterall, 2016; Schaich, 2019, for recent reviews).

[14] More technically, the origin of the problem comes from the so-called "spectrum doubling symmetry" of the naïve lattice action (see Montvay and Münster, 1994, sec. 4.4.1, for more detail). In a nutshell, the naïve discretization of a continuum QFT gives rise to *too much* symmetric structure, which must be broken in some way or another (or re-used for some other purpose) for the issue to disappear. Note as well that taking either the forward or the backward derivative for the fermionic field allows us to replace the

The fermion doubling problem might be thought to disappear at sufficiently low energies ($k \ll 1/a$) since the dispersion relations in the lattice and continuum models become approximately the same in this case, i.e., $E \sim k$ in the massless case. But this is misleading. Energy-momentum is conserved up to translation by $2\pi/a$ in the Brillouin zone. So two arbitrarily soft incoming fermions with momentum $k \sim 0$ can for instance give rise even at tree level to two different kinds of outgoing fermions with momentum $k \sim \pi/a$ and opposite chirality (e.g., Karsten and Smit, 1981, pp. 107-8; Smit, 1986, p. 4). Insofar as the lattice action involves interaction terms, we could thus observe in principle redundant fermions for arbitrarily low-energy inputs.

Now, as often in physics, we can solve the issue of interest by distorting the theory in some way or another. There are, in fact, many different kinds of solutions in this case. But they all come at a cost. The so-called "Wilson fermions" and "staggered fermions" solutions to the fermion doubling problem provide two popular examples. The first consists in adding a new non-renormalizable term to the action, which gives rise to an additional effective mass for fermion doublers (Wilson, 1977). This effective mass becomes arbitrarily large for a sufficiently small lattice spacing a, which means that the problematic fermion doublers become "invisible" at sufficiently low energies. The procedure, however, comes at the cost of chiral symmetry, and this for any value of the lattice spacing (see, e.g., Moore, 2003, pp. 32-3, for a short summary of the main issues affecting this solution).

The second solution consists in using the sixteen doublers in four dimensions to define four Dirac spinors and spread their components appropriately over each hypercube of the lattice—hence the label "staggered" (Kogut and Susskind, 1975; Banks et al., 1976; Susskind, 1977). These spinors may be reinterpreted in terms of a new kind of flavor and matched to traditional fermions in the continuum limit when each hypercube goes to a point. This is sufficient to make the original issue disappear. But for any non-zero value of the lattice spacing, the original flavor and translational symmetries of the model are lost (see, e.g., Capitani, 2003, sec. 7, for a discussion).

As it happens, the Nielsen-Ninomiya theorem provides principled reasons to believe that there is no fully satisfactory solution (Nielsen and Ninomiya, 1981a,b,c; Friedan, 1982). At a heuristic level, this theorem states that it is impossible to put chiral fermions on a lattice without introducing new fermionic particle species, breaking chiral symmetry, or violating some other desirable principle. A more precise yet still intuitive formulation in Euclidean space goes as follows. Suppose that the fermionic kinetic part of

problematic $\sin(ak)$ by $\sin(ak/2)$. But in this case, the action is not Hermitian (see, e.g., Capitani, 2003, p. 46, for more detail).

a lattice action takes the following form (for a lattice of infinite extent):

$$S[a] = \frac{1}{a}\int_{\mathrm{B_Z}} \frac{d^4k}{(2\pi)^4}\bar{\psi}(-k)D(k)\psi(k), \tag{12}$$

where the integral runs over the first continuous Brillouin zone ($\mathrm{B_Z}$), $\bar{\psi}$ and ψ refer to fermionic field variables in momentum space, and $D(k)$ stands for the kinetic operator in momentum space. Then, the theorem states that the following four conditions cannot be all satisfied: (i) $D(k)$ is continuous in the Brillouin zone, which is equivalent to assuming that the kinetic operator in position space is local, in the sense that it decreases exponentially with distance; (ii) $D(k)$ reduces to the continuum fermionic operator in the continuum limit, i.e., $D(k) \sim \gamma^\mu k_\mu$ for $k \ll 1/a$; (iii) $D^{-1}(k)$ has only a pole at $k = 0$ and thus no fermion doublers; (iv) $D(k)$ preserves chiral symmetry, i.e., $\{\gamma_5, D(k)\} = 0$ (see, e.g., Moore, 2003, sec. 8, for pedagogical proofs).[15]

The scope of the Nielsen-Ninomiya theorem is also far-reaching. In the Euclidean form stated above, it applies to any regularization involving a partially compact momentum space with periodic boundary conditions. We may of course obtain the very same unwanted consequences by regularizing the model in a different way. For instance, the standard implementation of dimensional regularization for chiral theories breaks chiral invariance (see, e.g., Collins, 1986, sec. 13.2, for more detail). But in contrast to lattice models, this kind of intermediary violation need not be a defining feature of the model. In particular, we may perfectly well define a continuum EFT by integrating out some high-energy field configurations, compute relevant perturbative quantities with the help of dimensional regularization methods, and remove the regulator at the end of the procedure to obtain chirally invariant quantities. We can also explicitly delineate the domain of the continuum EFT through boundary conditions on its renormalized correlation functions.[16]

[15] One might wonder about the relevance of this theorem since chiral invariance is broken in the SM. The issue reappears in this case through the appearance of incorrect corrections to the chiral current. For instance, in the simple case where we have fermion doublers, the contribution is such that the chiral anomaly cancels exactly since we have the same number of left-handed and right-handed fermions. But having the correct chiral anomaly is crucial to obtaining correct predictions, say, about the decay of neutral pions.

[16] Relatedly, Nielsen and Ninomiya's claim that the theorem extends to any regularization scheme seems to be rather heuristic (1981a, p. 222). They merely point to the failure of familiar regularization schemes like the Pauli-Villars and dimensional regularization schemes to satisfy all the conditions of (a slightly different version of) the theorem.

5 Discussion

I have argued so far that lattice QFTs are too different from their effective continuum counterparts to serve as their foundational proxies. We can certainly use lattice QFTs (in perturbation theory) to obtain good enough numerical estimates for the dynamics and correlation functions of continuum EFTs. But their underlying mathematical structure and the physical systems they depict are far from alike even at low energies. In particular, for any non-zero lattice spacing, the dynamics and correlation functions associated with each system involve different physical degrees of freedom and a number of different core principles and symmetries. The Nielsen-Ninomiya theorem even provides principled reasons to believe that a latticization introduces an irreducible incompatibility between lattice QFTs and continuum EFTs.

There are two kinds of responses one might have in mind at this point: (i) resist the claim that the latticization introduces significant physical differences at low energies; (ii) grant the existence of such differences but take lattice QFTs to provide a new foundation for continuum EFTs. As I will now argue, response (i) is not compelling and response (ii) faces a serious obstacle.

Consider response (i) first. RG-informed discussions of realistic effective QFTs suggest a natural way out (e.g., Wallace, 2006; J. D. Fraser, 2018; Williams, 2019). We might insist that significant physical deviations between lattice QFTs and continuum EFTs, be it through their background space, degrees of freedom or symmetries, are ultimately all parametrized by irrelevant contributions to the action, correlation functions, and predictions that become increasingly negligible at low energies. And we might take this to support the claim that lattice QFTs are ultimately "scale-l equivalent" to continuum EFTs, in the sense that they generate approximately isomorphic algebras of operators at sufficiently low energies relative to a reference cut-off l (Wallace, 2006, p. 48), or, in the path integral formalism, that their correlation functions are approximately isomorphic to each other in this regime (Wallace, 2011, p. 122).

Although RG-informed strategies are largely successful in my view when applied to continuum EFTs and lattice QFTs separately, the situation is more complicated when we compare them with each other. Let me pick two particularly salient cases discussed above to illustrate this point.

First, although the contributions of dynamical lattice artifacts originating from the lack of spatiotemporal constraints become increasingly small at low energies, the lattice system remains to a large extent as much spatiotemporally non-covariant at low and high energies. Granted: lattice translations do become increasingly refined for arbitrarily small a. So I agree with Wallace (2006, p. 51) that the lack of translation covariance is a "small-scale property" of lattice QFTs that we can ultimately deem as physically insignificant at low energies. But there is no sense in which the energy-momentum conservation law becomes increasingly similar to its continuum counterpart at low energies: we can in principle produce arbitrarily high-energy outputs from low-energy inputs even in the case of an arbitrarily large first Brillouin zone $k \in]-\pi/a, \pi/a]$ in the limit $a \to 0$. There is also no sense in which rotations of lattice structures by $\pi/2$ become "infinitesimal" when we decrease a. There is likewise no sense in which irreducible representations of the discrete rotation group become increasingly similar to their continuum counterparts in this limit. And the same can of course be said of the causal and local structure of the lattice system, whether we speak of its metric or non-local interactions. In a word: many built-in lattice properties hold as much at low- and high-energy scales.

Second, the Nielsen-Ninomiya theorem also holds indifferently across scales. For a start, the most popular solutions to the fermion doubling problem do not merely confine its effects to the irrelevant dynamical sector of lattice models. Take staggered fermions for instance. Strictly speaking, they do not give rise to relevant dynamical artifacts. But the staggered lattice action still violates the spin-flavor structure of its effective continuum counterparts at arbitrarily low energies (among other issues). The case of Wilson fermions is even more straightforward: we obtain both relevant and irrelevant contributions upon renormalization due to the loss of chiral invariance. We may safely neglect irrelevant contributions at low energies. But for relevant contributions like a quark mass renormalization, we need to fine-tune the bare parameters of the Wilson action as we move toward low energies (e.g., Capitani, 2003, p. 44). And this, in turn, arguably signals a physically significant impact of the lattice at low energies, in close analogy with the naturalness problem in the Higgs case (Williams, 2015).

But perhaps we should not be overly worried about eliminating remnant low-energy lattice artifacts "by hand". Since there is no reason to interpret the lattice overly strictly as some real discrete structure lying at short distances, we should arguably not take overly seriously remnant artifacts of the latticization at long distances. Fine-tuning the bare parameters of a lattice action merely amounts in this case to compensating for an overly drastic idealization.[17]

There are two main things to say in response here. (i) Relevant dynamical artifacts are not tied to specific lattice models. Rather, they enjoy a certain degree of universality, in the sense that they can be generated by any kind of regularization or any kind of matching to a high-energy theory that violates the low-energy constraint of interest (e.g., chiral invariance). So even if we do not take any particular latticization or any specific solution to the fermion

[17] I am thankful to David Wallace for pressing me on this point.

doubling problem too seriously, there is still room for being worried about the physical significance of the relevant dynamical terms they may generate. (ii) Interpreting this specific kind of fine-tuning procedure in formal terms does not seem to be perfectly neutral with respect to the content of high-energy physics. For all we know, the next high-energy theory ahead may well break chiral invariance as crudely as Wilson fermions do, and there may well be a high-energy mechanism that prevents the appearance of relevant chiral-symmetry-breaking contributions at low energies. In this case, we would be justified in reinterpreting the fine-tuning procedure as a model-independent implementation of this mechanism at low energies.

Now, physicists may ultimately find a lattice solution to the fermion doubling problem that fully confines its effects to that of irrelevant dynamical terms. Yet this prospect seems unlikely for two reasons.

First, the four conditions in the Euclidean version of the Nielsen-Ninomiya theorem outlined above cannot be violated only at high energies, strictly speaking. Regarding condition (i), the continuity of $D(k)$ is a global property of the kinematic structure of fermionic fields that holds as much at low and high energies. Condition (ii) is even worse: it may be read as a strictly low-energy condition, i.e., $D(k)$ reduces to the continuum fermionic operator at low energies $k \ll 1/a$. Giving up condition (iii) is no better: as we saw above, we could in principle produce high-energy doublers with low-energy incoming fermions. And violating condition (iv) opens up the door for relevant contributions otherwise forbidden by chiral invariance.

Second, as far as I can tell, all the most well-known solutions to the fermion doubling problem like SLAC fermions and Ginsparg-Wilson fermions (including domain wall, perfect and overlap fermions) do irreducibly present some remnant mark of it at arbitrarily low energies. For instance, Ginsparg-Wilson fermions break the standard chiral symmetry for arbitrary a directly at the level of the standard kinetic term (which is a relevant term). So although the departure from standard chiral invariance becomes increasingly negligible in the continuum limit, i.e. $[\gamma_5, D] = aD\gamma_5 D$, the fact that it is broken is still carried by the kinetic properties of arbitrarily low-energy fermions (see, e.g., Tong, 2018, sec. 4.4, for an introductory discussion).[18]

We might follow two lines of retreat at this stage. The first is to content ourselves with the existence of a continuum limit for lattice QFTs. They may be too mathematically and physically different from their effective continuum counterparts at any non-zero energy scale ak. But this does not

[18] Relatedly, I am sympathetic to the overall spirit—but not the specifics—of Ruetsche's remark: "Although the details of [mirror fermions'] appearance vary with the details of the lattice spacing and model, the *fact* of their appearance does not." (2020, p. 311) Strictly speaking, there are many ways to get rid of fermion doublers *simpliciter*. So their appearance is far from factual. But whichever solution we choose, it always comes at a cost at any finite energy.

undermine the existence of a smooth connection between them at $a = 0$ and thus our ability to track how lattice QFTs provide a mathematically well-defined reformulation of continuum EFTs, one might say.

Yet there are good reasons to doubt the success of this line of retreat. For a start, we do not yet have a well-defined and non-trivial continuum limit of realistic lattice QFTs. Any attempt to ground continuum EFTs with their lattice counterpart in this way is thus blocked in the very first step. But there is worse. As already emphasized, if realistic lattice QFTs had a well-defined and non-trivial continuum limit, they would immediately lose their foundational relevance. We would be able to provide a clear foundation for continuum EFTs out of their "internal" non-perturbative UV completion, regardless of whether we may come up one day with a well-defined and empirically successful "external" non-perturbative UV completion (presumably in the form of a theory of quantum gravity).

The second line of retreat is to adopt a more austere interpretation of realistic effective QFTs in terms of the numerical information encoded across a background space by low-energy correlation functions. On this view, we should for instance not attribute any kind of physical significance to the degrees of freedom, symmetries, and other higher-order principles at play in any given model. Their physical content must be entirely expressible numerically by means of correlation functions across arbitrary space-time points. And in this case, it does seem that lattice QFTs and their effective continuum counterparts have approximately the same physical content at sufficiently low energies. We indeed do not need to take seriously correlations over the continuum—looking at sufficiently coarse-grained regions is enough. Irrelevant interaction terms also bring negligible numerical contributions to correlation functions. And this interpretative move allows us to simply ignore the physical underpinning of ad hoc fine-tuning procedures and lattice-dependent deviations from existing principles (as for space-time and chiral symmetries).

I see two main issues with this line of retreat. (i) On this view, lattice QFTs do not really provide a foundation for their effective continuum counterparts but rather replace them altogether. All that matters is indeed to recover the approximate numerical value of correlation functions across some background space. There is thus no need to account for the background mathematical structures, physical content and principles of continuum EFTs, or even explain their inner workings and successes. (ii) In practice, lattice QFTs and their effective continuum counterparts are meant to obtain numerical information on largely non-overlapping domains. Although lattice perturbative techniques have been developed to fill the gap and help to compare their respective predictions, lattice QFTs tend to be highly impractical in the perturbative regime. Continuum EFTs, by contrast, are typically used in sufficiently well-behaved perturbative regimes. So, at least in practice, there is no sense in which lattice QFTs are meant to be used as a means to derive approximately the same correlation functions as their effective continuum counterparts.

This brings us at last to the opposite reaction, i.e., response (ii) above. We may decide to embrace the multifaceted mathematical and physical differences existing between lattice QFTs and their effective continuum counterparts and take the former to provide a new definition for realistic effective QFTs. On this view, continuum EFTs are downgraded to a set of efficient perturbative computational schemes with no proper foundation. The advent of lattice QFT even undermines the need to provide any such foundation.

But this response faces a severe obstacle, which is already familiar in the philosophical literature on axiomatic approaches to the foundations of QFTs. In Ruetsche's (2011, p. 11) and Williams' (2019, p. 2) words, lattice QFTs are far from "discharging" realistic effective QFTs of their "scientific duties," whether we speak of prediction or explanation. As already emphasized, using lattice QFTs in the perturbative regime is indeed highly impractical, especially when it comes to computing higher-order loop contributions. And it is not as if lattice practitioners have been able (or even willing) to use lattice QFTs to reproduce the entire set of perturbative predictions that underwrite the remarkable success of their effective continuum counterparts.

I should emphasize that the empirical and explanatory standing of lattice QFTs is still in much better shape than that of their algebraic counterparts. With enough fiddling and tuning, we can formulate any kind of realistic scalar, fermionic, and gauge model in four dimensions. Lattice QFTs also come with improved non-perturbative explanations and predictions, whether we speak of confinement or hadron masses for instance. Yet there are still both practical and principled limitations to lattice QFTs' ability to supplant their effective continuum counterparts. I have already mentioned their limited perturbative power. But it is worth noting that despite remarkable progress during the last two decades, lattice numerical simulations are still severely limited in terms of computational cost (see Aoki et al., 2025, for the latest FLAG review). And on the theoretical side, as we have already seen, using lattice QFT for models involving fermions typically involves roundabout solutions with intrinsic limitations. For instance, although staggered fermions constitute the computationally most effective method, it only enables us to account for four fermionic species, which is of course unsuitable for the six existing quark flavors.

Before concluding, let me emphasize that there may well be various ways of steering a middle course, acknowledging both the radical departure of lattice QFTs from their effective continuum counterparts and their inability to replace them altogether. We might for instance think of lattice QFTs as "foundational band-aids:" they allow us to assign a clear mathematical and conceptual meaning to some defective component parts of a target theory and thus clarify their physical meaning independently of its remaining component parts. The path integral measure provides a particularly clear example: lattice QFT suggests that a continuum measure merely corresponds to a well-defined continuum product of well-known measures at each space-time point. But although this avenue is worth exploring, it is likely to face its own issues too. In particular, it is not entirely clear to me whether we can interpret the central component parts of a theory in complete isolation. As we saw above, discretizing the path integral measure typically generates new dynamical artifacts. And this suggests that we cannot interpret the path integral measure of a continuum EFT completely independently of its dynamics (and inversely).

6 Conclusion

I have argued that the middle-path strategy of using lattice quantum field theory to provide a proper mathematical and conceptual foundation for continuum effective field theories faces a two-sided obstacle. The most interesting side in my sense is theoretical: realistic lattice QFTs are just too different from realistic continuum EFTs even at low energies to serve as their foundational proxies. In particular, we have seen how replacing their continuum background space-time had drastic higher-order physical implications at the level of the dynamics and correlation functions. Perhaps most surprising of all is the stubborn resistance of lattice QFT to accommodate fermionic fields in the standard way. Changing gears and taking lattice QFTs to provide an altogether new foundation for their effective continuum counterparts does not work either. We face again a severe albeit more straightforward obstacle in this case: lattice QFTs are far from reproducing all their empirical and explanatory successes.

What should we do then? Although I do not have the space to develop and defend this here, let me conclude by suggesting two lessons. First, we should grant that lattice QFTs play their own heuristic, explanatory, and computational roles independently of continuum EFTs, be it the search for a non-perturbative UV-complete QFT or the computation of low-energy non-perturbative quantities. That is, lattice QFT is probably best seen as a self-standing framework in the multi-faceted toolkit physicists use to study QFT (as opposed to a substitute framework developed to give a proper foundation to continuum EFTs). Second, we should endorse a form of methodological pluralism about standards of mathematical rigor, acknowledging the well-foundedness of more or less formal interpretative projects. In particular, although we may not have the clearest mathematical sense of how effective continuum field configurations are interlocked with one another, continuum EFTs are still presumably sufficiently well-defined to address at least some significant foundational and conceptual issues. If we follow this route, however, the challenge will be to understand how, exactly, we are supposed to integrate the various insights gained from those interpretative projects together.

Acknowledgments: I am very grateful to John Dougherty, Benjamin Feintzeig, James Fraser, Michael Miller, Daniele Oriti, David Wallace, audiences in Paris and Pittsburgh, and two anonymous referees, for very fruitful exchanges and helpful feedback. This work was supported by the European Research Council [grant no. 101161899].

References

Aoki, Y. et al. (2025). "FLAG Review 2024". URL: `http://arxiv.org/abs/2411.04268`.

Bain, J. (2013). "Effective Field Theories". In *The Oxford Handbook of Philosophy of Physics*. Ed. by R. Batterman. Oxford: Oxford University Press, pp. 224–254.

Baker, D. J. (2016). "The Philosophy of Quantum Field Theory". *Oxford Handbooks Online*. URL: `http://philsci-archive.pitt.edu/11375/`.

Banks, T., Susskind, L., and Kogut, J. (1976). "Strong-coupling Calculations of Lattice Gauge Theories: (1 + 1)-dimensional Exercises". *Physical Review D* 13.4, pp. 1043–1053.

Beneke, M. (1999). "Renormalons". *Physics Reports* 317.1, pp. 1–142.

Bergner, G. and Catterall, S. (2016). "Supersymmetry on the Lattice". *International Journal of Modern Physics A* 31.22, p. 1643005.

Blum, A. (forthcoming). *Probing the Consistency of Quantum Field Theory*. 2 vols. Cambridge: Cambridge University Press.

Butterfield, J. and Bouatta, N. (2015). "Renormalization for Philosophers". In *Metaphysics in Contemporary Physics*. Ed. by T. Bigaj and C. Wuthrich. Leiden: Brill Rodopi, pp. 437–485.

Capitani, S. (2003). "Lattice Perturbation Theory". *Physics Reports* 382.3, pp. 113–302.

Collins, J. (1986). *Renormalization: An Introduction to Renormalization, the Renormalization Group and the Operator-Product Expansion*. Cambridge: Cambridge University Press.

Creutz, M. (1985). *Quarks, Gluons and Lattices*. Cambridge: Cambridge University Press.

Degrand, T. A. and Detar, C. (2006). *Lattice Methods for Quantum Chromodynamics*. World Scientific Publishing Company.

Djukanovic, D., Gegelia, J., and Meißner, U.-G. (2018). "Triviality of Quantum Electrodynamics Revisited". *Communications in Theoretical Physics* 69.3, p. 263.

Dougherty, J. (2021). “I Ain’t Afraid of No Ghost”. *Studies in History and Philosophy of Science* 88, pp. 70–84.

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

Earman, J. (2020). “Quantum Physics in Non-Separable Hilbert Spaces”. URL: `http://philsci-archive.pitt.edu/18363/`.

Earman, J. and Fraser, D. (2006). “Haag’s Theorem and its Implications for the Foundations of Quantum Field Theory”. *Erkenntnis* 64.3, p. 305.

Fischer, J. (1997). “The Use of Power Expansions in Quantum Field Theory”. *International Journal of Modern Physics A* 12.21, pp. 3625–3663.

Fletcher, D. (forthcoming). “Similarity Structure on Scientific Theories”. In *Topological Philosophy*. Ed. by B. Skowron. URL: `http://philsci-archive.pitt.edu/17222/`.

Fraser, D. (2009). “Quantum Field Theory: Underdetermination, Inconsistency, and Idealization”. *Philosophy of Science* 76.4, pp. 536–567.

— (2011). “How to Take Particle Physics Seriously: A Further Defence of Axiomatic Quantum Field Theory”. *Studies in History and Philosophy of Modern Physics* 42.2, pp. 126–135.

— (2020). “The Development of Renormalization Group Methods for Particle Physics: Formal Analogies Between Classical Statistical Mechanics and Quantum Field Theory”. *Synthese* 197.7, pp. 3027–3063.

Fraser, J. D. (2018). “Renormalization and the Formulation of Scientific Realism”. *Philosophy of Science* 85.5, pp. 1164–1175.

— (2020a). “The Real Problem with Perturbative Quantum Field Theory”. *The British Journal for the Philosophy of Science* 71.2, pp. 391–413.

— (2020b). “Towards a Realist View of Quantum Field Theory”. In *Scientific Realism and the Quantum*. Ed. by S. French and J. Saatsi. Oxford: Oxford University Press, pp. 276–292.

Fraser, J. D. and Rejzner, K. (2024). “Perturbative Expansions and the Foundations of Quantum Field Theory”. *The European Physical Journal H* 49.1, p. 10.

Friedan, D. (1982). “A Proof of the Nielsen-Ninomiya Theorem”. *Communications in Mathematical Physics* 85.4, pp. 481–490.

Georgi, H. (1993). “Effective Field Theory”. *Annual Review of Nuclear and Particle Science* 43.1, pp. 209–252.

Hunt, J. (2024). “Staying On-shell: Manifest Properties and Reformulations in Particle Physics”. *Synthese* 204.4, p. 116.

— (forthcoming). “On the Value of Reformulating”. *Journal of Philosophy* ().

Isidori, G., Wilsch, F., and Wyler, D. (2024). “The Standard Model Effective Field Theory At Work”. *Reviews of Modern Physics* 96.1, p. 015006.

Itzykson, C. and Drouffe, J.-M. (1989). *Statistical Field Theory: From Brownian Motion to Renormalization and Lattice Gauge Theory*. Vol. 1. Cambridge: Cambridge University Press.

Karsten, L. H. and Smit, J. (1981). "Lattice Fermions: Species Doubling, Chiral Invariance and the Triangle Anomaly". *Nuclear Physics B* 183.1, pp. 103–140.

Kogut, J. and Susskind, L. (1975). "Hamiltonian Formulation of Wilson's Lattice Gauge Theories". *Physical Review D* 11.2, pp. 395–408.

Langacker, P. (2017). *The Standard Model and Beyond*. 2nd edition. London: Taylor & Francis Inc.

Luke, M., Manohar, A. V., and Savage, M. J. (1995). "Renormalons in Effective Field Theories". *Physical Review D* 51.9, pp. 4924–4933.

Maas, A. (2020). "Lattice Quantum Field Theory". URL: `https://static.uni-graz.at/fileadmin/%5C_Persoenliche%5C_Webseite/maas%5C_axel/lattice2020.pdf`.

Miller, M. E. (2018). "Haag's Theorem, Apparent Inconsistency, and the Empirical Adequacy of Quantum Field Theory". *The British Journal for the Philosophy of Science* 69.3, pp. 801–820.

— (2021). "Worldly Imprecision". *Philosophical Studies* 178.9, pp. 2895–2911.

— (2023). "Mathematical Structure and Empirical Content". *The British Journal for the Philosophy of Science* 74.2, pp. 511–532.

Mitsch, C., Gilton, M., and Freeborn, D. (2024). "How Haag-Tied is QFT, Really?" *Philosophy of Physics* 2.1.

Montvay, I. and Münster, G. (1994). *Quantum Fields on a Lattice*. Cambridge: Cambridge University Press.

Moore, G. D. (2003). "Informal Lectures on Lattice Gauge Theory". URL: `https://theorie.ikp.physik.tu-darmstadt.de/qcd/moore/latt%5C_lectures.pdf`.

Nachtergaele, B., Sims, R., and Young, A. (2019). "Quasi-locality Bounds for Quantum Lattice Systems. I. Lieb-Robinson Bounds, Quasi-local Maps, and Spectral Flow Automorphisms". *Journal of Mathematical Physics* 60.6, p. 061101.

Nielsen, H. B. and Ninomiya, M. (1981a). "A No-go Theorem for Regularizing Chiral Fermions". *Physics Letters B* 105.2, pp. 219–223.

— (1981b). "Absence of Neutrinos on a Lattice: (I). Proof by Homotopy Theory". *Nuclear Physics B* 185.1, pp. 20–40.

— (1981c). "Absence of Neutrinos on a Lattice: (II). Intuitive Topological Proof". *Nuclear Physics B* 193.1, pp. 173–194.

Rivat, S. (2019). "Renormalization Scrutinized". *Studies in History and Philosophy of Modern Physics* 68, pp. 23–39.

— (2021). "Effective Theories and Infinite Idealizations: A Challenge for Scientific Realism". *Synthese* 198.12, pp. 12107–12136.

Rivat, S. (2025). “Open Systems Across Scales”. *Synthese* 205.1, p. 11.
— (forthcoming). “Wait, Why Gauge?” *The British Journal for the Philosophy of Science* ().
Rosaler, J. (2015). “Local Reduction in Physics”. *Studies in History and Philosophy of Modern Physics* 50, pp. 54–69.
Rosaler, J. and Harlander, R. (2019). “Naturalness, Wilsonian Renormalization, and ”Fundamental Parameters” in Quantum Field Theory”. *Studies in History and Philosophy of Modern Physics* 66, pp. 118–134.
Ruetsche, L. (2011). *Interpreting Quantum Theories: The Art of the Possible*. Oxford: Oxford University Press.
— (2020). “Perturbing Realism”. In *Scientific Realism and the Quantum*. Ed. by S. French and J. Saatsi. Oxford: Oxford University Pres, pp. 293–314.
Schaich, D. (2019). “Progress and Prospects of Lattice Supersymmetry”. In *Proceedings of the 36th Annual International Symposium on Lattice Field Theory*. Vol. 334. SISSA Medialab, p. 5.
Schwartz, M. (2013). *Quantum Field Theory and the Standard Model*. Cambridge: Cambridge University Press.
Smit, J. (1986). “Fermions on a Lattice”. *Acta Physica Polonica B* 17, p. 531.
— (2002). *Introduction to Quantum Fields on a Lattice*. Cambridge: Cambridge University Press.
Susskind, L. (1977). “Lattice Fermions”. *Physical Review D* 16.10, pp. 3031–3039.
Tong, D. (2018). “Gauge Theory”. URL: `https://www.damtp.cam.ac.uk/user/tong/gaugetheory/gt.pdf`.
Wallace, D. (2006). “In Defence of Naivete: The Conceptual Status of Lagrangian Quantum Field Theory”. *Synthese* 151.1, pp. 33–80.
— (2011). “Taking Particle Physics Seriously: A Critique of the Algebraic Approach to Quantum Field Theory”. *Studies in History and Philosophy of Modern Physics* 42.2, pp. 116–125.
Williams, P. (2015). “Naturalness, the Autonomy of Scales, and the 125GeV Higgs”. *Studies in History and Philosophy of Modern Physics* 51, pp. 82–96.
— (2019). “Scientific Realism Made Effective”. *The British Journal for the Philosophy of Science* 70.1, pp. 209–237.
Wilson, K. G. (1977). “Quarks and Strings on a Lattice”. In *New Phenomena in Subnuclear Physics: Part A*. Ed. by A. Zichichi. The Subnuclear Series. Boston: Springer, pp. 69–142.