The fate of particles in finite-temperature quantum field theory

Adam Koberinski[1*]

[1*]Rotman Institute of Philosopby, University of Western Ontario, 1151 Richmond Street, London, N6A 3K7, Ontario, Canada.

Corresponding author(s). E-mail(s): akoberin@uwo.ca;

Abstract

In this paper I provide a new context in which one should be skeptical of the applicability of one form of particle ontology for quantum field theory. I show that the approximate, non-fundamental particle ontology that conventionally underwrites scattering theory in the vacuum fails to apply to finite-temperature quantum field theory. Thus, the scope of an emergent particle ontology is further restricted for the framework of quantum field theory as a whole. I discuss the implications of this failure for understanding particlelike phenomenology in finite-temperature applications of the Standard Model, most notably the quark-gluon plasma. While some particle phenomenology can be recovered, several further features typically associated with particles are lost in this domain.

Keywords: Quantum field theory, Ontology, finite-temperature field theory

1 Introduction

Of great foundational interest in the philosophy of physics is the question of fundamental ontology for quantum field theory (QFT). While the broader question of ontology appropriate to quantum theories is typically focused on non-relativistic theories, the relativistic domain of QFT introduces a host of new foundational challenges. The standard approach to fundamental ontology of QFT takes some form of a field ontology, where the fields are thought of as wave functionals (Sebens 2022), or as referenced to states of spacetime regions (Swanson 2020, 2024). The field ontologies are usually presented as a clear contrast to an alternative family of particle ontologies (cf. (Tumulka 2022, Sec. 6.5.2) and Dürr and Lazarovici (2020)). While the two options do not exhaust the space of logical possibility, it is fair to characterize the debate as a debate between particle versus field ontology. A standard background assumption is that there is one true (or best) fundamental ontology suitable to the QFT framework.[1]

There are clear motivations for both camps. Motivating the particle ontology is the fact that it is a natural reading of the practice of particle physics, where QFTs are most heavily used. The paradigmatic experiments in particle physics involve colliding beams of subatomic particles together, and analyzing the resulting products of collision. If QFT does not admit a particle ontology, the argument goes, then it is difficult to make sense of the practice of particle physics, and hence the main source of empirical support for QFTs. Motivating the field ontology is the basic idea that the framework of QFT is a *field-theoretic* framework. QFTs are constructed by starting with a theory of classical fields, then quantizing. It therefore seems most plausible that the underlying ontology would be a field ontology. While the motivations are easy to state, the difficulties lie in developing a positive account of either ontology. Even if we set aside the measurement problem, it seems that the QFT formalism offers several barriers to elaborating either option in a traditionally satisfactory way. First, empirically successful QFTs are plagued with mathematical ambiguities and/or inconsistencies, making it difficult to even get clear on the purported representational content of the theory and its resulting ontology. Second, quantum fields are far removed from the objects of everyday experience, making their possible ontological underpinnings challenging to intuit. They combine the difficulties of classical fields with quantum indeterminacy, and compound both with a fundamentally relativistic causal structure. The standard strategy has instead been to provide arguments *against* a particular (class of) ontology(ies). This project is somewhat easier: one generates a list of properties characteristic of the ontology in question, and argues that those features cannot generically hold within the QFT framework.

The idea of a fundamental particle ontology, in particular, has been the subject of severe criticism by philosophers and physicists over the last several decades (Fraser 2021). On the basis of classical intuitions, we typically require particles to have definite, finite extent; to have some *haecceity* (primitive identity); to be countable; aggregable; and be the bearers of some definite properties, such as charge and mass. Moving to relativistic physics changes the last condition: we expect the particles to obey the relativistic mass-energy relationship $E = \sqrt{p^2 + m^2}$, but otherwise the classical expectations are retained. As I will elaborate in Sec. 2, nearly all of these properties cannot hold in general in QFT, while still others fail to hold for the realistic interacting QFTs that form the basis of the Standard Model of particle physics. However, some concession is still needed to connect QFT to particle physics, where it boasts some of the most precise predictions in the history of science. If a *fundamental* particle ontology is untenable, then something else must allow for and explain the phenomenological *appearance* of particles in particle physics experiments. One orthodox candidate for an approximate, non-fundamental particle ontology can be recovered in the form of the interaction picture and the scattering matrix (S-matrix). This requires a notion of a 'particles' applicable at asymptotic times before and after the scattering event, and the associated Fock space that comes with it. This is approximate and non-fundamental

[1] For the purposes of this paper, I use the term 'QFT' to refer to relativistic QFTs.

because it only approximately holds in the specific context of particle scattering, and may fail to apply elsewhere. While this is already a significant change from the project of fundamental ontology, and introduces further mathematical ambiguities, it forms the folk ontology of particle physics that most physicists subscribe to. If tenable, it suffices for explaining the success of scattering experiments in particle physics and the presence of particles as detected in these experiments.

In this paper I argue that the standard particle ontology underlying scattering experiments in the vacuum sector of QFT cannot be extended to the context of finite-temperature QFT. Assuming, for the sake of argument, that the approximate, non-fundamental particle ontology used in the S-matrix construction is generally suitable to purpose, the goal of this paper is to further circumscribe it as applicable only to situations in which temperature effects can be neglected. The argument is rather simple: the approximate, non-fundamental particle ontology used in scattering theory depends on having a well-defined notion of a state of definite particle number, and approximating particles as non-interacting far before or after a scattering interaction. In the context of finite-temperature QFT, both of these requirements become mathematically and physically unreasonable. The thermal background state means that finite-temperature effects are unitarily inequivalent to any vacuum Fock space, blocking the use of number operator eigenstates as definite particle states. Further, persistent thermal fluctuations and interactions spoil the idea of approximately free particles at asymptotic times. As a result, one cannot define the S-matrix at finite temperature. Any particle ontology that depends on the S-matrix, therefore, is untenable in this context.

Finite-temperature QFT is the application of QFT for systems where thermal effects cannot be neglected. It involves using a background thermal state at some temperature $T = 1/k\beta$, rather than the zero-temperature vacuum state. It differs from standard applications of condensed matter physics in that it treats relativistic systems, typically those described by the Standard Model of particle physics. As I will discuss below, the early universe and quark gluon plasmas are two places where finite-temperature effects become relevant for QFTs. While various formal features of condensed matter physics find analogues in finite-temperature QFT, there are important differences as well. A secondary goal of this paper is to bring philosophical attention to finite-temperature QFT more generally.

The main result does not impact the empirical contact between scattering experiments and particles in vacuum QFT, and therefore has no impact on standard scattering experiments in particle accelerators. It does, however, show the further limits of this particle ontology for QFT. Even if we accept that a particle ontology is at best approximate and non-fundamental, one conclusion of this paper it that it must emerge in a context-dependent way. When we move from the vacuum sector to the thermal sector of a given QFT, this ontology cannot hold even approximately. What we have instead are limited calculable quantities that are traditionally interpreted as aspects of a particle phenomenology. While these suffice for understanding thermodynamical aspects of QFT, and may underwrite a *different* particle notion even further removed from a unified foundation in theory-level mathematical structures, they do not meet the demands of underwriting the phenomenological appearance of something particlelike. One appropriate general lesson to draw is that articulating an ontology must be done in a context-specific manner, and that we cannot expect any non-fundamental ontology to apply to the QFT framework as a whole.[2] At best, only negative ontological conclusions can be drawn at the fully general framework level of QFT. This general point is similar in spirit to recent criticisms of *pristine* interpretations of QFT (Dougherty 2023; Ruetsche 2011, 2024; Ruiz de Olano et al. 2022). This result is different in kind from other temperature-related effects in QFT, like the Unruh effect, though both can be leveraged to provide arguments against some form of particle ontology. Arageorgis et al. (2003), for example, use the Unruh effect to illustrate the non-uniqueness of a particle interpretation in the case of a free, vacuum QFT. My argument begins with QFT defined over a thermal state, and in general factors in interactions.

In Sec. 2 I recap some earlier results against a particle ontology, and articulate the resulting approximate, non-fundamental particle ontology that is nevertheless invoked as an explanation for particle behaviour in scattering experiments. Sec. 3 contains the main result, arguing that this weakened particle notion is inapplicable to finite-temperature QFT. In Sec. 3.1 I discuss the prospects of a particle ontology in asympotitically free QFTs, such as quantum chromodynamics, and pose several problems for applying the vacuum scattering ontology to this domain. Sec. 4 discusses more minimal, operational particle notions for QFT, suitable to the finite-temperature sector. Some combination of a generalized LSZ formula and a detector-based approach might suffice for recovering particle phenomenology in the finite-temperature contexts where particles seem to arise. I conclude that these provide resources to explain some aspects of particle phenomenology, but amount to something even further removed from traditional particle ontology than is found in the vacuum sector for particle scattering. Sec. 5 draws broader conclusions for the ontology of QFT.

2 Particle ontology in vacuum QFT

There are several results that undermine the prospects for a fundamental particle ontology for QFT. The argument strategies all take the following form: take a set of properties to be constitutive of the particle concept. Then, show that those properties cannot be maintained in the QFT framework. For example, several authors have argued that there cannot be *localizble* particles in QFT, in the sense of being definitely contained in any finite region (Malament 1995; Halvorson 2001; Halvorson and Clifton 2002). It is a well-known result that particle number in QFT is not conserved; particles can be created and destroyed. Finally, even in the context of non-relativistic quantum mechanics, identical particles lose their haecceity, and this carries over to QFT. This has led to a proposal for a weakened notion of particles, that Teller (1997) has called 'quanta'. The defining features of quanta are that they carry the discrete properties associated with classical particles (charge, mass-energy relations,etc.) as well as

[2] Presumably, the goal of a fundamental ontology is to be universally applicable across contexts within QFT. While there are several results indicating that QFT is not hospitable to a fundamental particle ontology, there are far less convincing reasons to be skeptical of a fundamental field ontology. One might argue that the project of articulating any universally applicable fundamental ontology for QFT should fail, as Ruetsche (2011) contends. It is outside the scope of this paper to defend or attack the prospects of a fundamental ontology for QFT.

new quantum properties (spin, colour, etc.) and are aggregable: there exist states of a definite number of quanta, and combining a state containing n quanta with a state containing m quanta gives rise to a state with $n+m$ quanta. But even this weakened quanta concept is severely limited: Fraser (2008) argues that it is only applicable to the special case of non-interacting QFTs, at least if one wants a fundamental particle ontology. Standard conventional physics wisdom, however, uses this free quanta ontology as the basis for scattering theory via the S-matrix and LSZ reduction formula. Bain (2000); Wallace (2001) defend different forms of approximate or emergent particle ontology on the basis of scattering theory, and it is a plausible demand that some form of particle ontology must emerge to make sense of the practice of particle physics, even if it is only approximate and non-fundamental. While one may question whether recovery of a particlelike phenomenology for scattering experiments deserves the title of "ontology", I set that issue aside.[3] Instead, I articulate the most common form of this non-fundamental particle ontology here, starting with the free quanta picture, then incorporating the interaction picture for S-matrix theory, and finally introducing a sketch of the LSZ reduction formula.[4]

Free quanta are found in the formalism of noninteracting QFT via the Fock space representation. Take, as an example, a free Bosonic real scalar field with positive mass, given in a particular frame by $\phi(\mathbf{x},t)$, or covariantly by $\phi(x)$.[5] The symmetries of Minkowski spacetime allow for a Fourier transform of this field, split into positive and negative frequency components:

$$\phi(\mathbf{x},t) = \int \frac{d^3k}{(2\pi)^{3/2}\sqrt{2\omega_k}} \left(a^\dagger(\mathbf{k},t)e^{ik\cdot x} + a(\mathbf{k},t)e^{-ik\cdot x}\right), \tag{1}$$

with $\omega_k^2 = k_0^2 + m^2$. The $a^\dagger(\mathbf{k},t)$ and $a(\mathbf{k},t)$ operators can be interpreted as operators that create and annihilate free quanta of momentum $\mathbf{k}$ at time t, respectively. The reason why depends on the structure of the Fock space. Part of the definition of the Fock space is that there exists a unique vacuum state $|\Omega\rangle$—a state with no particles. This state has the property that

$$a(\mathbf{k},t)\,|\Omega\rangle = 0 \qquad \text{for all } \mathbf{k}. \tag{2}$$

Free Bosonic fields with any finite number of particles can be written in the Fock space representation, for which the Hilbert space $\mathcal{F}(\mathcal{H})$ is a direct sum of the n-fold symmetric tensor products of the associated one-particle Hilbert space $\mathcal{H}$:

[3] While the in and out states characterizing S-matrix elements are discussed here as connected to particle phenomenology, this is still very far from actual experimental detection and inference to outgoing particles from a scattering event. These reconstructions are complex tasks laden with expectations on the basis of theory (Beauchemin 2017). I take for granted here that these processes genuinely license a demand for finding 'particles' somewhere within QFTs. I thank an anonymous reviewer for pressing this point.

[4] If the use of the term 'ontology' in this relatively loose form is objectionable, the reader can replace it with their preferred term applicable at this level of approximation. What is important for the discussion that follows is that this emergent ontology has an exactly specifiable, consistent formal expression at the foundational level of QFT.

[5] The construction here generalizes to nonscalar fields, either Fermionic or Bosonic. Subtlety is needed in working with massless fields (cf Wald (1994); Fraser (2008)).

$\mathcal{F}(\mathcal{H}) = \bigoplus_{n=0}^{\infty}(\bigotimes^n \mathcal{H})$.[6] The one-particle Hilbert space is spanned by a basis of vectors generated from $|\Omega\rangle$ by a single application of $a^\dagger(\mathbf{k}, t)$ for all possible $\mathbf{k}$ and t (see Wald (1994) for further details).

Within a Fock space representation, it is possible to define a number operator $N(\mathbf{k}, t) \equiv a^\dagger(\mathbf{k}, t)a(\mathbf{k}, t)$, whose eigenstates take the form:

$$N(\mathbf{k}, t)\left[\frac{a^\dagger(\mathbf{k}, t)^n}{\sqrt{n!}}\,|\Omega\rangle\right] = n\left[\frac{a^\dagger(\mathbf{k}, t)^n}{\sqrt{n!}}\,|\Omega\rangle\right] \equiv n\,|n(\mathbf{k})\rangle\,, \tag{3}$$

and can be interpreted as states of n free quanta of momentum $\mathbf{k}$ at time t. From this number operator, the total number operator at time t, $N(t) = \int d^3k N(\mathbf{k}, t)$ is also well-defined as the sum of the individual number operators for each k.

What makes the Fock space a good host for quanta is the following. First, the total number operator provides a count of how many quanta there are at any given time, fulfilling the aggregability criterion. When acting on the vacuum state, the total number operator gives zero. Second, this Fock space representation is relativistically covariant, as required for a particle interpretation of relativistic QFT. Third, the Lorentz-covariant Fock representation on Minkowski spacetime for a given field with given mass is unique up to unitary equivalence; there is therefore a single, well-defined particle notion for a given field.[7] Finally, the n particle states have the correct energy-momentum relations for corresponding relativistic particles. The Hamiltonian for a free Bosonic field in the Fock representation is given by $H = \int d^3k\, \omega_k N(\mathbf{k})$, such that an n quanta state is an eigenstate of the Hamiltonian with eigenvalue $\omega_{k1} + \omega_{k2} + \cdots + \omega_{kn}$, which corresponds to the classical energy-momentum relations for relativistic particles. When restricted to free QFT, quanta form the basis of a perfectly acceptable particle ontology, albeit one with fewer 'particlelike' properties than we may have first expected.

As Fraser (2008) argues, however, this quanta representation for free fields cannot be carried over to interacting fields directly. Haag's theorem shows that Hilbert spaces for free and interacting particles are unitarily inequivalent. This means that free quanta cannot directly fill the role of quanta for realistic interacting theories. Fraser further argues that possible candidate particle representations dealing directly with the interacting theory require one to give up more and more of the features inherent in even the weakened quanta ontology. Fraser concludes that interacting QFT therefore does not "support the inclusion of particlelike entities as fundamental entities in our ontology" (p. 856). The prospects for a *fundamental* particle ontology for interacting QFTs therefore seems dim.

Nevertheless, in practice, physicists continue to use the free Fock space representation as the basis for a particle ontology for weakly interacting theories. The resulting

[6] For Fermionic fields, take the antisymmetric tensor product instead.

[7] On might think that the Unruh effect is an exception to this uniqueness. In the Unruh effect, accelerating observers may take themselves to be confined to the right Rindler wedge, which is both a subset of Minkowski spacetime and a globally hyperbolic spacetime in its own right. The appropriate particle concept for the Rindler spacetime is inequivalent to that of Minkowski spacetime, but this is due to the different symmetry structures of the Minkowski spacetime and the Rindler wedge spacetime. Since the Rindler observer is eternally accelerating, the spacetime has a different timelike Killing vector, leading to a different decomposition of the field into positive and negative frequency modes.

particle concept is non-fundamental, and the mathematical representation is inexact, but it serves to connect the formalism of QFT to the phenomenology of particle physics. The standard approach to doing so relies on a notion of asymptotically free particles as employed in the S-matrix, which is used to calculate the probability of a particle collision involving some specified set of particles (the "in" state) resulting in a different set of specified particles (the "out" state). The physical idea in simple: elements of the S-matrix are simply the amplitudes for incoming state $|a\rangle_{in}$ transitioning to outgoing state $|b\rangle_{out}$, where we take the limit that the incoming state is defined in the asymptotic past, and the outgoing state in the asymptotic future:

$$S_{ba} = \lim_{t\pm \to \pm\infty} \langle b|_{out} \, U(t_+, t_-) \, |a\rangle_{in} \, , \tag{4}$$

where $U(t_+, t_-)$ is the unitary time evolution operator. At this point, the states of particles in the asymptotic past and future are taken to be isolated enough to be treated as approximately free, justifying the use of free states of definite quanta. In introductory textbook presentations, this is accomplished via the interaction picture, which is physically plausible only in weakly coupled theories.

In the interaction picture, one splits the Hamiltonian H into a free part and an interacting part: $H = H_0 + H_I$, such that the evolution under H_0 is analytically soluble, and H_I can be treated as a small perturbation to the free dynamics. The interaction picture is a hybrid of the Schrödinger picture—under which states evolve in time while operators remain constant—and the Heisenberg picture—under which operators evolve in time while the states remain constant. For the interaction picture, operators take the form:

$$O_{IP}(t) \equiv e^{iH_0 t} O_S e^{-iH_0 t}, \tag{5}$$

while states take the form:

$$|a; t\rangle_{IP} \equiv e^{-iH_I t} \, |a\rangle_H \, , \tag{6}$$

where the subscripts IP, S and H denote interaction picture, Schrödinger picture, and Heisenberg picture, respectively. In the interaction picture, states are subject to evolution under the interaction Hamiltonian H_I, while operators evolve under the free Hamiltonian H_0. Then, the interaction unitary is determined as a series expansion in the interaction part of the Hamiltonian:

$$U(t_+, t_-) = \sum_{n=0}^{\infty} \frac{(-i)^n}{n!} \int_{t_-}^{t_+} dt_1 dt_2 \ldots dt_n \mathrm{T}\{H_{I,IP}(t_1) \cdots H_{I,IP}(t_n)\}, \tag{7}$$

where $\mathrm{T}\{\cdot\}$ is a time ordering operator (which puts the operators with the earliest time arguments to the left and those with later arguments to the right), and

$$H_{I,IP}(t) \equiv e^{iH_0 t} H_I e^{-iH_0 t} \tag{8}$$

is just the interaction picture form of the interaction Hamiltonian. The S-matrix can then take incoming and outgoing states to be states of definite number and types of free quanta, and evolve these with the full evolution operator $U(t_+, t_-)$ to find the transition amplitude. It is the interaction picture that exploits the free quanta notion, and uses these free quanta as asymptotic states. But this is at best a non-fundamental particle ontology, since the treatment of particles as asymptotically free is an idealization. Interactions are always present, even at asymptotic times. The idealization is considered acceptable in the context where the asymptotic particles are far apart, since interaction strength dies off with distance.

Beyond the asymptotically free idealization, this ontology is complicated by the fact that, strictly speaking, the interaction picture is mathematically ill-defined for QFT. It presupposes that the states for the free and fully interacting theories can be formulated within the same Hilbert space, and this is proven to be impossible due to Haag's theorem (Earman and Fraser 2006; Fraser 2008; Koberinski 2023). Haag's theorem roughly states that, for relativistic QFTs with infinite degrees of freedom, a free theory is unitarily inequivalent to an interacting theory. So we cannot use free field quanta states and evolve them with an interacting unitary evolution operator, since this operation is ill-defined in QFT. The idealization from the interaction picture is therefore known to fail.

There are two ways to respond. First, is to argue that Haag's theorem does not apply in actual applications of the interaction picture. Second is to develop a more exact, but conceptually weaker, particle concept that applies more generally. Start with the former; in practice, physicists often evade Haag's theorem by regularizing the theory, which is often done by introducing an upper and lower length limit for domains of integration. This is a necessary step in extracting well-defined predictions from an interacting QFT, and it also serves to evade the assumptions of Haag's theorem by rendering the QFT finite, in the sense of having a finite number of degrees of freedom (Duncan 2012; Miller 2018). This solves the issue for making sure calculations are well-defined, but renders the resulting physical interpretation opaque. We have the interaction picture idealization that justifies taking free quanta as input and output states in the S-matrix, but only when we further alter the QFT formalism to make the field degrees of freedom finite.[8] Typically, remornalization is needed in addition to regularization, such that regularized quantities only appear at intermediate stages of the calculation. Thus the connection between free quanta and particles present in scattering experiments is circuitous. However, if we set aside scruples regarding mathematical rigour, we have some notion of approximate, non-fundamental particles for QFT, based on free field quanta and the interaction picture.

But what about the latter option? For vacuum field theories, there exists a more general prescription for determining S-matrix elements that does not rely on the interaction picture: the LSZ reduction formula (Peskin and Schroeder 1995). This provides an exact (i.e., nonperturbative) relationship between S-matrix elements and correlation functions, though one that is still dependent on the asymptotic properties of the field operators in the noninteracting limit. The dependence is mathematically weaker than unitary equivalence, automatically avoiding the assumptions needed for Haag's

[8]One additional problem with regularizing the QFT is that the resulting theory is no longer Lorentz covariant. There are other possible solutions to Haag's theorem that depend on different forms of scattering theory or weaker forms of equivalence between free and interacting theories (Earman and Fraser 2006; Koberinski 2023). Many of these are more physically transparent, principled evasions of Haag's theorem, but suffer the problem of not being directly applicable to calculation in realistic interacting QFTs. Perhaps the most well-known, Haag-Ruelle scattering theory also suffers similar issues to the LSZ formalism, discussed below. In all cases, the connection between free quanta and the interacting theory is made more tenuous.

theorem. The downside is that the interpretation of the connection to free quanta is also weaker, leaving us with a particle concept that is more tenuously connected to free quanta. The requirement is that, for any pair of normalized states $|\psi_1\rangle$, $|\psi_2\rangle$ and any normalizable solution $f(x)$ to the free field equations, the interpolating field $\phi(x)$ yields all the same expectation values as the free field $\phi_0(x)$ in the limit asymptotic time limit:

$$\lim_{t\to-\infty}\int d^3x\,\langle\psi_2|\,f(x)\overleftrightarrow{\partial_t}\,\phi(x)\,|\psi_1\rangle = \int d^3x\,\langle\psi_2|\,f(x)\overleftrightarrow{\partial_t}\,\phi_0(x)\,|\psi_1\rangle\,, \tag{9}$$

with $f\overleftrightarrow{\partial_t}g = f\partial_t g - g\partial_t f$. A similar relation holds for the $t\to\infty$ limit. The interpolating field is defined within the interacting theory, and this relation implies that the expectation values using free fields and those using the interpolating fields agree in the asymptotic limit. So we are working with states that are empirically indistinguishable from free quanta states in the asymptotic past and future.

While S-matrix elements are important for connecting QFTs to empirical domains via scattering experiments, correlation functions arise more naturally from the theory itself. Connecting the two is therefore of fundamental importance. The n-point correlation functions are vacuum expectation values of time-ordered products of n field operators

$$G_n = \langle\Omega|\,\mathrm{T}\{\phi(x_1)\cdots\phi(x_n)\}\,|\Omega\rangle\,. \tag{10}$$

Working in momentum space, the Fourier transform of G_n is given by:

$$\Gamma(p_1,\cdots,p_n) = \int\prod_{i=1}^{n}\left[d^4x_i e^{ip_i\cdot x_i}\right]G_n. \tag{11}$$

The LSZ reduction formula for Klein-Gordan scalar particles[9] relates an S-matrix element consisting of m incoming particles of momentum $p_1,\cdots,p_m$ and $n-m$ outgoing particles of momentum $q_{m+1},\cdots,q_n$ to the n-point correlation function Γ, as a function of the interpolating fields:

$$\langle q_{m+1},\cdots,q_n|p_1,\cdots,p_m\rangle = \prod_{i=m+1}^{n}\left[\frac{i(q_i^2-m^2)}{2\pi^{3/2}Z^{1/2}}\right]\prod_{i=1}^{m}\left[\frac{i(p_i^2-m^2)}{2\pi^{3/2}Z^{1/2}}\right]\Gamma(q_{m+1},\cdots,q_n,-p_1,\cdots,p_m), \tag{12}$$

with Z a (re)normalization factor introduced in relating the free fields to the interacting fields. This formula says, roughly, that asymptotic particles in the S-matrix are given by the residues of poles in Γ as four-momenta are constrained to be on-shell.[10] While the LSZ formula is far more general, it is still limited to cases of massive, unbound particle states. In the strong interaction limit, bound states become more

[9] By this I mean scalar field theories whose free part obeys the free Klein-Gordan equation, plus interaction terms. The LSZ reduction formula is more general, but becomes even more cumbersome to write for more general fields. See any textbook treatment, like Peskin and Schroeder (1995) for further details.

[10] We will see in Section 4 that this particular feature might still generalize to the finite-temperature case, though the underlying free quanta concept might not.

prominent, altering the mass spectrum and severing the connection with asymptotically free particle states. However, generalizations can be formulated by working directly with bound state operators, recovering some notion of particle phenomenology.

While more useful in practice, the LSZ particle concept is even more tenuously connected with standard expectations of particlelike properties. The remaining connection is still parasitic on the free quanta concept. While the relation is no longer a perturbative relation in the strength of the coupling, it is a weak convergence of norms between the interacting field and the free field at asymptotic times. Thus, the LSZ particle states are *not* free quanta states, but are operationally indistinguishable in the context of asymptotic scattering. The interpretation of the S-matrix in and out states are still conceptually dependent on free quanta, despite them not being truly present in the formalism (though, see Bain (2000) for a subtly different interpretation). Whether one sticks to the interaction picture or uses the LSZ formalism, the vacuum interacting quanta concept is parasitic on the vacuum free quanta concept in the asymptotic limit.

Recall that the free quanta concept in QFT is more limited than an ordinary notion of classical particles. These free quanta are not localizable and have no haecceity. However, despite the fact that free quanta cannot be exactly localized, they can be *approximately* localized in a relativistically covariant manner. Roughly speaking, the overlap between approximately localized, free relativistic one-particle states will drop off exponentially with distance. So free quanta are at least approximately localizable (Halvorson and Clifton 2002; Colosi and Rovelli 2008). To the extent that we accept a non-fundamental particle ontology for interacting QFT, we have some license to infer that scattering particles inherit this approximate localizability from the free quanta. This suffices for recovering phenomenology of particle detection in collision experiments. Thus, when using the free quanta for defining interacting particles, layers upon layers of approximations are introduced, but these bottom out in some well-defined formal representation in the theory.

To summarize, while a fundamental particle ontology faces several severe challenges, one can recover the link to particle physics with a non-fundamental particlelike ontology, where we have a well-enough defined notion of "approximately free, approximately localizable" quanta that fill the role of particles in particle physics. Setting aside mathematical scruples concerning the exactness of the relation between interacting fields and asymptotically free quanta, we have entities that approximate many of the desirable particle properties one might expect to hold. These are particle surrogates well-suited to cases where the S-matrix is used to model particle collisions with asymptotically free incoming and outgoing particles. It is often the case in other domains of physics that we have mathematically well-defined concepts that hold *exactly* of the theoretical structures, and that we then take these structures to approximately represent features of the world. What we have for particles in QFT instead are theoretical structures that themselves only *approximately* instantiate the well-defined mathematical concepts. These theoretical features are still expected to only approximately represent the actual features of the world. As we will see in the next section, there are other applications of QFT where even this highly approximate particle notion cannot be defined.

3 No asymptotic particles at finite temperature...

Now we are at a place to introduce finite-temperature QFT, and to argue that the approximate, non-fundamental particle interpretation from the last section cannot get off the ground in this context. The point here is a modest one: the general consensus is that the scattering particle ontology for QFT is at best an approximate, non-fundamental ontology that arises in certain special contexts. The argument here is that finite-temperature QFT is *not* one of those contexts. I will also provide some motivation for the fact that this particle ontology is not needed for many applications at finite-temperatures. Much of the discussion here illustrates only the basics of finite-temperature QFT; further details can be found in several textbook treatments. See, e.g., Le Bellac (2000); Zinn-Justin (2000); Blaizot et al. (2004); Kapusta and Gale (2007).

To start, note that QFTs at finite-temperature are typically only useful in the presence of interactions. Intuitively, the thermal background must interact with the states of interest in order to be dynamically relevant.[11] However, we can still define a free QFT at finite temperature, by simply replacing the vacuum state for the Fock space with a thermal state ρ_β at temperature $\beta = 1/kT$. This is a free theory still, since the equations of motion do not include interactions. One might hope that, if a particle notion exists in the finite-temperature free case, then one can construct an approximate, non-fundamental particle notion in full analogy with vacuum QFT. At the conceptual level, finite-temperature QFT is not that different from vacuum QFT. The only real change is the background state on the basis of which other states are defined. The same QFT governed by the same dynamical equations admits of representations based on a vacuum state and thermal states at different temperatures. The thermal states are mixed states, unlike the vacuum state, which is a pure state. Further, a choice of thermal state fixes a reference frame, as temperature is not a relativistically invariant quantity. Otherwise, the fundamental quantities of interest are still n-point functions, interpreted as time-ordered expectation values for field operators in the relevant state:

$$G_{n,\beta} = \mathrm{Tr}(\rho_\beta \mathrm{T}\{\phi(x_1)\cdots\phi(x_n)\}), \tag{13}$$

where the $G_{n,\beta}$ now also include an index to represent the particular temperature at which they are calculated. Recall that, for the vacuum case, asymptotically free quanta were well-defined, and interacting theories depended on this free quanta concept for particle content. Dealing directly with the S-matrix, the interaction picture provided one way to connect free and interacting quanta. At the level of the n-point correlation functions, one could generate the link to the S-matrix via the LSZ reduction formula. In either case, we need a free particle concept, and then to connect this to interacting theories via the S-matrix. Unfortunately, there is no S-matrix for finite-temperature QFT.

[11] There are some caveats here. In applications of quantum statistical mechanics, for example, noninteracting QED at finite temperature provides a foundation for understanding blackbody radiation. Further, quantum chromodynamics (QCD) in the high-temperature limit is asymptotically free. The strong couplings therefore approach zero in the limit of high-energies/high-temperatures. QCD will be discussed further below.

The reason behind this is quite simple. Take the definition of the S-matrix from the previous section. It relies on the idealization that incoming and outgoing particle states are approximately non-interacting excitations to the vacuum state, so that the free Fock space representation is appropriate for cataloguing the input and output content of a scattering experiment. A Fock space is constructed from the vacuum state by repeated application of the creation operators for the appropriate quanta, allowing for states with a definite particle number to be properly aggregable. But at finite-temperatures, the background field is not the vacuum: it is a thermal state at some non-negligible temperature. This change introduces two problems. First, representations built from a thermal state are unitarily inequivalent to vacuum representations, so one cannot directly use the free vacuum representation as a basis for the free thermal representation. While there are at least satisfactory workarounds to this unitary inequivalence in the vacuum sector, like the LSZ formalism, there are no similar tricks in the thermal case.[12] This means that there are no states of definite vacuum-particle number in finite-temperature QFT, where a definite number state is taken to be an eigenstate of the appropriate number operator. Naively taking the expectation value of the total number operator (defined in Sec. 2) on any thermal state, we get a value of infinity. Start by writing the appropriate thermal state in a Fock basis. The thermal state is then a mixed state containing a superposition of each of the n-particle states, weighted by the Gibbs factor:

$$\rho_\beta = \sum_{n=0}^{\infty} \int d^3k \frac{1}{Z} e^{-\beta E_n} |n(\mathbf{k})\rangle \langle n(\mathbf{k})| \,, \tag{14}$$

where Z is the partition function, included to ensure that ρ_β is normalized, $\mathrm{Tr}(\rho_\beta) = 1$. Taking the expectation of the total number operator yields infinity, though the number operator for a finite region yields a finite expectation value. Strictly speaking, however, unitary inequivalence implies that the thermal state is not in the domain of the total number operator, so a different representation is needed.

Note that it is not just a problem with the thermal state itself: *no states* in the thermal representation will be even approximate eigenstates of the number operator. Therefore no states in a thermal representation of a QFT will have a well-defined particle number: the thermal state is not an eigenstate of the Hamiltonian or number operator. Without an accompanying solution to the measurement problem, it is therefore incorrect to say that the thermal state is *actually* a state with a definite number of particles, even if that number is infinite. Even ignoring the problem of unitary inequivalence, we still arrive at the conclusion that the thermal state has an indefinite number of particles, as do all states created by applying the appropriate creation operators to a thermal state ρ_β. Within this formalism, one cannot determine the particle

[12] Such a picture would have to have interactions *and* temperature fall to zero as time approaches $\pm\infty$. As far as I am aware, no such picture is used in physics applications. Note the distinction here with the Unruh effect, in which the vacuum for free quanta Fock space appropriate to the Rindler wedge appears as a thermal state in the full Minkowski vacuum state. This is an inequivalence between vacuum representations on different spacetime backgrounds, not an inequivalence between different state choices on the same background spacetime.

content of any states created from a thermal state, even in the case of a noninteracting Hamiltonian. Since all states in finite-temperature QFT are constructed from the thermal state, the analogous construction of a particle notion for free fields fails.

Within the context of algebraic QFT, a stronger result can be shown. Chaiken (1968) shows that any representation of the Weyl canonical commutation relations that contains a number operator must be unitarily equivalent to a combination of Fock representations. Since the thermal representation is unitarily inequivalent to a Fock representation, it cannot contain a number operator.[13] But what about an analogue of a Fock space, suitable for a thermal background? While this would be an inequivalent particle notion to that of the vacuum Fock space, it might suffice to provide a basis for a particlelike ontology in the finite-temperature sector. However, there are good reasons to think that there are no generalizations of dynamically relevant particles in the thermal context. Narnhofer et al. (1983) show that, if quasi-particles exist with respect to a KMS state, then there is no particle scattering, i.e., the S-matrix is the identity. While Landsman (1988) is more optimistic about the possibility of a thermal analogue to vacuum perturbation theory, he acknowledges that "[t]he disparity between vacuum and thermal boundary conditions... is so large that in our opinion the standard perturbation scheme that copies the vacuum one is totally inadequate" (p. 143). In this case, it's not at all clear how one would relate the particle ontologies in the vacuum and at finite-temperature. Would we have to conclude that heating up the background induces a change in the ontology of the theory, even in the absence of phase transitions?

There are some equivalences between vacuum and thermal particle concepts in special cases, but these serve to illustrate the challenge of transferring interpretation from one representation to another.[14] First are the well-known cases like the Unruh effect and Hawking radiation, which relate acceleration in a vacuum spacetime to a thermal state. One important thing to note for the Unruh effect is that accelerating observers confined to the Rindler wedge can define their own vacuum and Fock space of Rindler particles, based on the spacetime symmetries of the Rindler wedge. These are unitarily inequivalent to the particles defined with respect to the Minkowski vacuum; it is with respect to the latter that an accelerating observer measures a thermal bath. Particle content is thus dependent on the spacetime structure in which it is embedded. Similarly for Hawking radiation, which relates quantum states on a curved spacetime to an asymptotically flat Minkowski particle notion. While there are no particles in the curved spacetime, the states look like a thermal bath with respect to the Minkowski vacuum. The important point is that a particle ontology needs to be suited to the *appropriate particle phenomenology*; if there is a mismatch, then the mere existence of a Fock-like representation will not suffice to legitimize particle talk.

Mathematical translation maps between a thermal representation of one theory and a vacuum representation of another are not sufficient to underwrite an interpretational equivalence or a transfer of ontological commitments from one to the other. We can see this with a second example. Jäkel and Robl (2014) establish a relativistically invariant version of the KMS condition used to pick out thermal states for a given QFT, and

[13] I thank an anonymous reviewer for the reference to Chaiken.

[14] Many of the results discussed in the following paragraphs are rigorously proved in the setting of algebraic QFT. For a thorough primer on algebraic QFT, see Ruetsche (2011).

prove that a toy model (the $\mathcal{P}(\phi)_2$-model) has all n-point functions satisfying this relativistic KMS condition. This is a construction of a thermal representation for the model. In doing so, they demonstrate an equivalence between the Schwinger functions for the thermal $\mathcal{P}(\phi)_2$-model on 2D Minkowski space at temperature β^{-1} and those for the vacuum $\mathcal{P}(\phi)_2$-model in a 2D Einstein universe of spatial circumference β. For the latter, the appropriate representation on which the Schwinger functions are defined is equivalent to the Fock representation. Given this equivalence, one could argue that the thermal $\mathcal{P}(\phi)_2$-model on Minkowski space has a suitable particle formalism, parasitic on the vacuum Einstein universe representation. One could just translate the thermal Schwinger functions to the Einstein universe vacuum, and read the particle content off of that representation. Again, however, these models differ in their spacetime structure. The particle content of the Einstein spacetime cannot be carried over to the Minkowski setting, as the spacetime setting is essential to defining a cogent particle concept. Further, the vacuum model is a closed universe with zero temperature, while the thermal model is an open universe at temperature β. There is no thermal particle concept by moving to the Einstein universe since that move also necessitates a move to zero temperature. This case raises intriguing questions about dualities in QFT that cannot be addressed here, but on the assumption that we have to provide a consistent physical interpretation to each of the two equivalent models as a whole, there is no clear avenue for borrowing the particle concept from one to apply to the other. Thus, the mere presence of *some* mathematically definable particlelike formalism is not sufficient for recovering a particlelike ontology suitable to the context.

Without a notion of asymptotic states of definite particle number, the S-matrix construction cannot get off the ground. But the problem gets qualitatively worse when we factor in interactions. The second problem is that, even if there were states of definite particle number to form the basis of the S-matrix, the motivation for treating particles in a thermal state as free at asymptotic times is undercut. The physical motivation behind using the free Fock space for a weakly coupled theory is that, far away from the collision interaction, the particles have spread out to the point of effectively no longer interacting with each other. This relies on local interactions and a background vacuum state, so that there is nothing persisting in the background for the far-flung particles to interact with. When the background state is a thermal state, however, this interpretation breaks down. Even if we had a well-defined notion of particles in this case, they would always be surrounded by a background thermal bath, and would have non-negligible interactions with the thermal excitations. So the physical justification for using the approximate, non-fundamental particle interpretation from the free vacuum QFT case is also undermined here, over and above any technical problems. It is no longer a plausible assumption that states are approximately free at asymptotic times.

So, we see that the strategy for recovering an approximate, non-fundamental particle ontology for vacuum QFT does not work in finite-temperature QFT. This implies that the scattering particle notion is even more circumscribed than is commonly discussed in the philosophical literature. But two lingering questions remain: do we recover any finite-temperature particle phenomenology? And if so, how? For many applications of finite-temperature QFT, a particle ontology is unnecessary. One important arena for finite-temperature effects is the early universe, where one is concerned with features of the fields like interaction strength (including mass) and symmetry breaking, and how these change as the universe expands and cools. The paradigmatic use for particle phenomenology—scattering theory—does not exist for finite-temperature QFT in this context. But other particlelike properties are important here. While the scattering picture is not used, we still talk of particles in many contexts in finite-temperature QFT. These applications range from stellar interiors, to nucleosynthesis, to high-temperature QCD. QCD will be discussed next, while more general questions of finite-temperature particle phenomenology will be discussed in Sec. 4.

3.1 ...Except for QCD?

QCD is the theory of the strong nuclear interaction, which posits three generations of quarks as the fundamental fermions, whose interactions are mediated by massless gluons. It is a Yang-Mills theory with Lie group SU(3), and describes the force responsible for the emergence and stability of hadrons like the proton and neutron, as well as nuclear binding. For the purposes of this paper, the most interesting feature of QCD is that it has a property known as *asymptotic freedom*: despite the strong coupling between quarks and gluons at low energies leading to quark confinement, the coupling strength decreases at higher energies, and approaches zero in the limit of infinite energy. Given the standard association of high energies with high temperatures, physicists expect that QCD over a high-temperature thermal state is a weakly coupled theory; in the limit $T \to \infty$ it becomes a free theory. Thus, one might think that the previous concerns about interactions in a thermal background state might be avoided, and it may be possible to use the pieces of the approximate particle ontology from the scattering picture in finite-temperature QCD, in a similar way that it is used for deep inelastic scattering in the high-energy regime. Additionally, there exists a QCD phase transition from a confinement phase to a phase of quark-gluon plasma, where the latter phase is (at least colloquially) described as a phase where quarks and gluons are able to travel freely, unconfined. So not only might it be possible, it might indeed be necessary to provide a particle ontology in the context of quark-gluon plasma. While the physical motivation for a particle ontology in these regimes is stronger, I will argue that the scattering formalism does not supply an adequate basis for a particle ontology, even in QCD.

Start with the general idea of asymptotic freedom. The renormalization group flow for the strong coupling is negative, meaning that the coupling strength g decreases at higher and higher energies μ: $\beta(g) = \mu\frac{\partial g}{\partial \mu} < 0$. This can be realized when collisions occur between highly energetic particles, or when there is an environment with a high ambient energy, such as a high-temperature environment. As the name 'asymptotic freedom' implies, at sufficiently high energies the strong coupling goes to zero, resulting in a theory of free quarks and gluons. The property of asymptotic freedom for QCD was first discovered by Gross and Wilczek (1973), and independently by Politzer

(1974), and was used to explain the phenomenon of deep inelastic scattering discovered at SLAC (Bloom et al. 1969).[15] When hadron collisions occur at sufficiently high energies, the hadrons behave as a collection of free point particles, resulting in deep inelastic scattering. These are particle collisions that require a highly energetic conditions, so one might think that finite-temperature QCD would be required to model the interactions. However, deep-inelastic scattering is modelled within the vacuum sector of QCD; all that is needed is that the incoming particles have sufficiently high energy-momentum to reduce the effective coupling strength in the collision interaction region. One still works within a scattering picture with asymptotic in and out states being stable excitations to the vacuum—in this case, the in and out states are hadrons. Within the scattering region, the hadron is taken to be a collection of quarks, one (or more) of which collides with the other incoming particles. The background state does not need to provide the energy necessary to weaken the strong coupling since this energy is provided by the colliding hadrons. However, asymptotic freedom applies to the full high-energy sector of QCD, beyond the original phenomenon of deep inelastic scattering, and implies a regime in which both finite-temperature effects and particlelike behaviour coincide.

Due to the weakening of the strong coupling at higher energies, QCD exhibits a phase transition, moving from a quark confinement phase at low energies to a quark-gluon plasma phase at high energies (Satz 2011). The theoretical treatment of the transition region and of empirically accessible quark-gluon plasma energies is carried out via lattice gauge theory, which involves discretizing the theory by putting it on a lattice of finite extent. Calculations of properties in lattice QCD are typically carried out numerically, and fall well outside the scope of the scattering picture.[16] One reason for the necessity of lattice QCD near the phase transition is that it occurs at energies too low for the weak coupling approximation. It turns out that deconfinement occurs well before quarks and gluons become sufficiently weakly interacting (Engels et al. 1982). Even within lab-made quark-gluon plasma, above the phase transition, the strong coupling is still too strong to utilize perturbation theory. This is an application of QCD in which the scattering particle concept does not apply in this domain, so any particle phenomenology in this domain must rest on different formal features of QCD. Alternatively, the particle phenomenology might be recovered via weaker notions discussed in Sec. 4.

But in the very early universe, temperatures can be high enough to create a weakly interacting quark-gluon plasma. This is the most likely regime in which the approximate particle ontology from vacuum scattering theory might apply. In fact, due to the simplification of near asymptotic freedom, high-temperature quark-gluon plasma is comparatively easy to study analytically. Even here, however, there are important

[15] For more on the historical development of QCD and the evidence leading to its acceptance as the theory of strong interactions, see Cao (2010); Koberinski (2021).

[16] Whether one can construct a suitable particle ontology for lattice gauge theories is beyond the scope of this paper. In general, however, discretizing the theory alleviates the concerns involved with having an infinite number of degrees of freedom. However, the Stone-von Neumann theorem still does not apply to lattice QCD, given that the topology of the phase space is not $\mathbb{R}^{2n}$. There are exact, algebraic constructions suitable to the kinematics of Hamiltonian QCD (Kijowski and Rudolph 2005), even extended to the infinite volume limit (Grundling and Rudolph 2013), which may serve as a starting point to analyzing possible particle contents of these theories.

differences between high-temperature perturbation theory and analogous vacuum perturbation techniques. As discussed in Sec. 3, one major problem is the lack of a number operator for the thermal representation. Even if we take the fully non-interacting limit $T \to \infty$, due to the presence of thermal fluctuations, there are no states which are eigenstates of a number operator. Since we need to be in the high-temperature limit of QCD, we can't neglect these thermal fluctuations as small. In general they should swamp any contributions from the lightest quarks and all gluons, since the latter are massless (Blaizot et al. 2004). This means that, at high-temperatures, there is no definite fact about the number of quarks or gluons in any state or any region of space. Thus, adding a single quark or gluon to the thermal background via a creation operator will not lead to a state with one single particle, nor will it appreciably increase the probability of detecting only a single particle in that state.

Now consider the case of weakly interacting QCD for finite T. Importantly, there is still a principled difference between weakly interacting fields and free fields, and this difference is harder to idealize away when the background state is a thermal state at significantly high temperatures. Interactions in a thermal background state therefore introduce new problems not already dismissed in the vacuum case. Perturbation theory in this limit leads to several distinct, often nonanalytic contributions to bulk quantities due to the interactions of thermal modes (Blaizot et al. 2004). This is not simply a feature of QCD in particular, but holds for any thermal perturbation theory in QFT. In the vacuum case, a single set of modes with definite mass-energy relations can fill the role of a particle, and in the absence of interactions these modes will persist and retain the definite quantum numbers associated with a particle. For a single particle state in an interacting theory, the properties are different from the free properties due to the presence of virtual interactions, but these interacting properties will similarly persist. In a thermal background, granting that the free single particle state can be defined, the reintroduction of interactions causes larger disruption. Because the background states is far more active, the interactions with a single particle state make fluctuating changes to its properties, to the point of losing any clear sense of definite properties like mass. This makes it hard to consistently identify a state defining a single particle. Even setting aside the issues with defining a Fock space or a number operator in a thermal representation, interactions make defining even an appropriate single particle state challenging. Without any good formal analogues of the S-matrix particle states for high-temperature QCD, there is thus no analogous scattering particle ontology.

However, there are clearly some features of a quark gluon plasma that are, at least approximately, analogous to a collection of weakly interacting particles. In the limit where interaction strength goes to zero, many of the thermodynamic quantities describing bulk features of the plasma are strongly analogous to their ideal gas counterparts (Blaizot et al. 2004). Additionally, properties typically attributed to particles—like mass and colour charge—remain meaningful and undergo changes in the plasma state, even in the presence of nonzero interactions. These quantities are calculated from the n-point functions of the theory using some form of thermal perturbation theory. There is thus some truth to the idea that some particlelike phenomenology comes out of thermal perturbation theory for weakly interacting QFTs, even if there is not a clear state space for particles in this regime. What we should make of this particle phenomenology will be discussed in the next section.

4 Prospects for finite-temperature particle phenomenology

So far, we have seen that the finite-temperature sector of QFT is inhospitable to the non-fundamental particle ontology that seems to suffice for recovering particle phenomenology in scattering experiments in the vacuum sector of QFT. But it is not particularly troubling, as there is little pressure to recover a general particle ontology outside of the context of scattering experiments in particle physics. However, some specific theories might have other, context-specific ways to define a different particle concept. High-energy QCD, for example, is a domain where at least some particle phenomenology appears. Other contexts where we model the physics of finite-temperature relativistic particles include the interior of stars, where nuclear fusion is supposed to occur, and hot dense environment of the early universe, where quarks and gluons condense into baryons, and plasma condenses into light nucleii. In this section I will discuss the general phenomenological features that one might need in finite-temperature QFT, which of these actually indicate underlying particle phenomenology, and how these are recoverable without some analogue of a Fock space to provide a clear particle-based state space. What we are left with is an even more attenuated notion of particle, further removed from the approximate, non-fundamental notion appropriate for scattering in the vacuum.

One major context in which finite-temperature effects are essential is in understanding the physics of the very early universe. In the ΛCDM model of cosmology, the universe starts in an initially hot, dense state, and evolution forward in time is marked by an expansion of the spatial scale factor and an accompanying cooling. Important effects in the very early universe can be attributed to the thermodynamics of the Standard Model interactions.[17] For many purposes, an explicit finite-temperature formulation is not directly needed, as qualitative or order-of-magnitude calculations can be made by considering the energies at which certain interactions stop becoming viable, and mapping these to certain epochs in the history of early universe expansion. Some effects, however, require a full finite-temperature treatment. First and foremost are various phase transitions in the early universe, including the electroweak phase transition and the QCD phase transition from quark-gluon plasma to confinement. The latter was discussed in some detail in the previous section. The electroweak phase transition is the transition where the $SU(2) \times U(1)$ symmetry group partially breaks, resulting in massive weak bosons, a massive Higgs boson, and the massless photon. However, one does not need to recover any distinctly particlelike phenomenology in this transition. Instead, all that is needed is the field-theoretic description. Masses and symmetry breaking can be determined as properties of the fields, without any need to model particlelike excitations, using the effective action formalism (Koberinski 2024).

[17] For more detailed discussion of the methodology of early universe cosmology, see Smeenk (2019); Koberinski and Smeenk (2024). For a historically-focused discussion of the development of finite-temperature QFT and its application to the early universe, see Koberinski (2024).

Importantly, most of these calculations are also non-perturbative, in the sense that they do not proceed by an expansion in powers of the coupling constant. The effective action is exact, and approximation methods used to get at the full effective action need not use a weak coupling expansion.[18] There is thus no need for a robust particle ontology, and no use of the scattering formalism needed for the previously discussed particle ontology.

Other thermal aspects relevant to properties in the early universe, however, seem to display more particlelike phenomenology. Properties like dressed masses/charges and the mean free path in a thermal state are properties best described in terms of particles. But it is not at all clear that they require a particle ontology for the QFT in question. At least at the level of modelling, one can input a test particle into the background state to determine properties like a screened mass or mean free path. However, in order to make sense of this modelling strategy, it is natural to interpret the background state as analogous to a gas of particles. In fact, what is often done in finite-temperature QFT is to use the strong formal analogies between the partition function in QFT and in statistical mechanics to calculate thermodynamic and statistical properties of interest while completely bypassing a concrete modelling approach, and any explicit mention of particles (cf. Kapusta and Gale (2007)). However, such properties bring to mind some aspects of particle phenomenology; after all, a mean free path for a test particle implies that it has an average time/distance it travels unimpeded before *interacting with something*, an idea that brings to mind something particlelike. Unless this is simply a case of reading too much physical significance into the formal analogies between QFT and statistical mechanics (Fraser and Koberinski 2016; Fraser 2020), it would be desirable to have some principled theoretical justification from which to infer a particle ontology suited to these contexts. I discuss two potential options, each with some desirable features and some drawbacks.

First, one can take features of the LSZ reduction formula as a basis for finite temperature particle concept. Recall that the asymptotic particle states for the S-matrix were given by residues of poles in the n-point functions for the theory, when those n-point functions were forced to be on mass-shell. While the S-matrix is not defined at finite-temperatures, the n-point functions are. We could decide to treat the right hand side of Equation (12) as indicative of particle content in the finite-temperature context, without the direct connection to the S-matrix. There are several advantages to this approach. First is its continuity with the most general form of a vacuum particle notion used in physics practice. We can just treat the same quantities as encoding particle content at finite-temperature as we do at zero temperature. Next, it provides a highly general formalism for determining particle content. Insofar as we can calculate the n-point correlation functions and the residues of their poles, we can make claims about the particle content of the theory. The limitations of restriction to massive, unbound particles remains, but introduces no further complications beyond the vacuum case. The only major drawback to this approach is that, without the S-matrix, there is no clear connection to directly measurable particles.

[18]Some of the clever ways of approximating the effective action still involve a power series in some parameter, but in general this need not be in powers of the coupling constant.

This leads naturally into the second option for particle phenomenology: the detector approach (Grimmer et al. 2021; Gale and Zych 2023). On this approach, one couples a QFT to a (typically) non-relativistic particle detector. These can be as simple as a qubit—a two state quantum system with a ready state and detection state—or might include more realistic features of a detector. Importantly, the form of interaction between the detector and quantum field is constrained to be approximately localized. Then, one has a 'particle' whenever the detector makes a positive measurement. This approach to particle phenomenology can be summarized in a slogan: *a particle is whatever a particle detector detects*. There are several advantages to this approach, including flexibility, practicality, and direct operational connection to experimental practice. In principle, one could use a detector model on any background state, including finite-temperature states. If one adopts a particle notion suited to the vacuum representation, then a detector of this form will sporadically click detecting a (thermal) bath of so-called particles when coupled to a thermal state. This is analogous to the point made in Sec. 3, that a thermal state can be written as a weighted superposition of an infinite number of possible definite particle states. The overlap between a one-particle state and a thermal state will therefore be non-zero.

This approach was first developed in the context of a finite-temperature effect: the Unruh effect (Earman 2011), where an constantly accelerating observer moving through Minkowski spacetime will observe an apparently thermal state, whose temperature is proportional to their acceleration. An Unruh-DeWitt particle detector was used as a way to model whether and how an observer would actually detect that they were in a thermal state. Now finding use in broader contexts like relativistic quantum information theory, the detector approach is currently the most general approach that allows one to recover particle phenomenology in a wide range of contexts.[19]

While it has several advantages, one major disadvantage for the detector-model approach to particle phenomenology is that the inference to even an approximate, non-fundamental particlelike ontology is severely attenuated. Given the fact that one can rewrite *any* state of a quantum field in a form of Fock basis (as in Eq. (14)), there will in general be nonzero overlap between the vast majority of possible quantum field states and a given n particle state. This means that there is at least some probability of a particle detector detecting n particles in *any* field configuration, even ones in which there is no other convincing reason to think of as n-particle states. The detector approach may therefore be too permissive in assigning particle content. In particular, the connection is severed between a positive detection in a given state, and the state bearing the appropriate quantum numbers corresponding to the defining properties of the collection of detected particles. This was one of the main reasons for taking the asymptotic Fock states from the vacuum scattering picture to be appropriate quanta states, and the closest one could get to a particle ontology. While the detector approach allows for a much broader set of contexts in which one can recover aspects of particle phenomenology, it does so at the cost of casting too wide a net. In some

[19]This is one approach to a broader problem of developing a general measurement theory for QFT. Outside of the context of scattering theory, which employs asymptotic states, there is little consensus on how to best define a local measurement theory in the QFT context. While detector models form one approach (Gale and Zych 2023; Grimmer et al. 2021; Grimmer 2023), others are working to develop an algebraic setting for quantum measurements (Fewster and Verch 2024). See Papageorgiou and Fraser (2024) for a summary of approaches and a philosophical analysis of the latter approach.

contexts, detector modelling might be a useful approach, but at the general level it is too susceptible to false positives to underwrite a new, non-fundamental particle ontology. At the very least, significantly more work would be needed to link detector phenomenology to some form of particle ontology. One way forward in this direction is to combine the detector approach with the generalization of the LSZ reduction formula: in this case, the detector coupling would help provide the empirical link that disappears with the lack of an S-matrix at finite temperature.

While the concept of asymptotically free particles from vacuum QFT cannot generalize fully to finite-temperature QFT, a combination of a (local) detector approach with the LSZ reduction formula may provide a particle concept suited to purpose for most applications where particles may arise. We lose some of the nice properties that come along with the (free) Fock space representation, like definite particle states and aggregability. However, this might not be surprising. Several desirable features of particles are lost in the move from classical to quantum physics, from quantum physics to vacuum QFT, and now from vacuum QFT to finite-temperature QFT. We are only left with the constraint that particles satisfy the appropriate relativistic mass-energy relations, and that we can detect them as discrete clicks in a non-relativistically modelled detector. Is this enough to underwrite a particle ontology for finite-temperature QFT? It is at least enough to recover particle phenomenology; perhaps this is the best we can do in terms of recovering an approximate, non-fundamental particle ontology.

5 Conclusions

Over the past several decades, the status of a particle ontology in QFT has been reduced from a candidate for fundamental ontology to an approximate, non-fundamental ontology. However, the applications for which this ontology are applicable are essential for QFTs to make empirical contact with scattering experiments in particle physics. I have argued in this paper that even the standard approximate, non-fundamental particle ontology based on the free quanta concept can only work in vacuum QFT, if at all. When one is concerned with finite-temperature QFT, the connection to free quanta is severed, and so asymptotic particles do not generalize from the vacuum sector. The already limited particle ontology for QFT is thus even more severely limited. However, we can generalize the LSZ particle concept, though the connection to the S-matrix disappears. We are therefore left with an even weaker particle phenomenology that finds use in understanding thermodynamic properties of a QFT.

Since a fundamental particle ontology seems untenable, does this mean that the field ontology wins out? Not necessarily. While a field ontology seems plausible for QFT, the devil is in spelling out the details of what commitments that brings with it. Baker (2009) has argued that field ontologies utilizing wave functionals face the same problems as particle ontologies.[20] If one takes the effective field theory view seriously, there is also the live option that QFT cannot admit a fundamental ontology at all (Miller 2021; Dougherty 2023). Finally, the dichotomy between field and particle ontologies may be a false one, so care is needed. Positive arguments in favour of a

[20] But, see replies to this challenge from Sebens (2022); Swanson (2024).

particular ontology are needed, in addition to negative arguments ruling out other ontologies.

Part of the challenge in spelling out a positive ontology for QFT is the fact that it is a theoretical framework, rather than a theory in its own right (Koberinski 2019). Physical content is sparse at the framework level, allowing for many different applications in different contexts. That flexibility, which is a pragmatic feature of a framework, also makes defining a uniform ontology challenging. Wallace (2020) makes this point in the context of quantum theory more generally: some systems will have an appropriate particle ontology, others fields, and still others something different. One must specify a more concrete theoretical context to have enough physical content for positive ontological claims. Plausibly, such a context will be sensitive to the appropriate types of empirical access we have to that domain. This fact explains why, while negative ontological claims abound for QFT, positive ones are much harder to articulate. We typically don't expect a uniform ontology for all classical theories, and QFT should be no different. If one accepts this background view of QFT, then it is perhaps not so surprising that the scattering particle ontology is a poor fit for finite-temperatures.

Setting aside the effective view for now, some open questions remain regarding the relative fundamentality of different sectors of a given QFT. I will take this space to outline these questions, but leave them unanswered for future work. Suppose that we are interested in constructing an ontological hierarchy for QFTs. Is it possible to order the various sectors, such that one particular sector is taken as the most fundamental? Given that the approximate particle ontology from vacuum QFT does not apply to finite-temperature QFT, what lessons should we draw for the relative fundamentality of particles suited to scattering contexts? There are three possible options, each with plausible motivations. First, one could argue that the vacuum sector is more fundamental, since a finite temperature state is just a phenomenological description of some highly excited vacuum state. In principle, whatever ontology is best-suited to the vacuum sector would be a more fundamental base on which the finite-temperature ontology would depend. Next, one could argue that the finite-temperature sector is more fundamental, on the basis of the usual reductionist inference in particle physics that higher-energies reveal the more fundamental features of a QFT. One is thus required to move to finite-temperature representations to elucidate the more fundamental. As an example in applications, very early universe cosmology is typically taken to be a window into the most fundamental aspects of the universe; in this domain, one needs finite-temperature QFT. A third response is to deny that there is an order of relative fundamentality to the two sectors. Motivating this answer is the inability to prioritize one of the motivations for the first two options over the other. I am partial to this third option, but leave the question open.

Finally, one might wonder about the importance of the approximate particle ontology one can recover in the vacuum sector. For the purposes of this paper, I set aside the major issues with understanding the interaction picture, and assumed that one could apply its particle formalism to actual scattering experiments done in the lab. But suppose we reject this move, is there still any way to recover particle phenomenology for scattering experiments? The moves made to recover a weaker particle phenomenology in finite-temperature QFT may provide a template for doing so. If we give up on the requirement of having a well-defined mathematical representation of particle-like entities at a general level in the mathematical formulation of the theory, then we can recover many of the features of particle phenomenology using numerical simulations, lattice methods, and determining non-analyticities in the n-point functions of the full theory. By paying close attention to experimental practice, we could reconstruct the particle concept needed in the appropriate empirical contexts. While I have only argued against extending the scattering picture's particle ontology to finite temperatures, one might take the real lesson to be that one ought to give up on that picture as grounding particle phenomenology anywhere in QFT.

Acknowledgements

I am grateful to Benjamin Feintzeig, Doreen Fraser, Chip Sebens, and Chris Smeenk for helpful discussions and feedback on this paper, and to audiences at the 2024 PSA poster session for stimulating discussion. This work was supported by a Social Sciences and Humanities Research Council of Canada Insight Grant, for the project *The Structure of Effective Theories*.

Declarations

Availability of data and material

Data availability statement: not applicable

Competing interests

The author has no competing interests to declare

Funding

This work was supported by a Social Sciences and Humanities Research Council of Canada Insight Grant, for the project *The Structure of Effective Theories*.

Authors' contributions

All work conducted by Adam Koberinski, the sole author.

References

Arageorgis, A., J. Earman, and L. Ruetsche. 2003. Fulling non-uniqueness and the unruh effect: A primer on some aspects of quantum field theory. *Philosophy of Science 70*(1): 164–202 .

Bain, J. 2000. Against particle/field duality: Asymptotic particle states and interpolating fields in interacting qft (or: Who's afraid of haag's theorem?). *Erkenntnis 53*(3): 375–406 .

Baker, D.J. 2009. Against field interpretations of quantum field theory. *The British Journal for the Philosophy of Science 60*(3): 585–609 .

Beauchemin, P.H. 2017. Autopsy of measurements with the atlas detector at the lhc. *Synthese 194*(2): 275–312 .

Blaizot, J.P., E. Iancu, and A. Rebhan 2004. *Thermodynamics of the High-Temperature Quark–Gluon Plasma*, pp. 60–122.

Bloom, E.D., D.H. Coward, H. DeStaebler, J. Drees, G. Miller, L.W. Mo, R.E. Taylor, M. Breidenbach, J.I. Friedman, G.C. Hartmann, and H.W. Kendall. 1969. High-energy inelastic $e-p$ scattering at 6° and 10°. *Phys. Rev. Lett.* 23: 930–934. https://doi.org/10.1103/PhysRevLett.23.930 .

Cao, T.Y. 2010. *From current algebra to quantum chromodynamics: A case for structural realism.* Cambridge University Press.

Chaiken, J.M. 1968. Number operators for representations of the canonical commutation relations. *Communications in Mathematical Physics 8*(2): 164–184 .

Colosi, D. and C. Rovelli. 2008. What is a particle? *Classical and Quantum Gravity 26*(2): 025002. https://doi.org/10.1088/0264-9381/26/2/025002 .

Dougherty, J. 2023. Effective and selective realisms. pp. 1–24. https://doi.org/10.1086/724978 .

Duncan, A. 2012. *The conceptual framework of quantum field theory.* Oxford University Press.

Dürr, D. and D. Lazarovici. 2020. Understanding quantum mechanics. *Cham: Springer* .

Earman, J. 2011. The unruh effect for philosophers. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics 42*(2): 81–97 .

Earman, J. and D. Fraser. 2006. Haag's theorem and its implications for the foundations of quantum field theory. *Erkenntnis 64*(3): 305–344 .

Engels, J., F. Karsch, H. Satz, and I. Montvay. 1982. Gauge field thermodynamics for the su (2) yang-mills system. *Nuclear Physics B 205*(4): 545–577 .

Fewster, C. and R. Verch 2024. *Measurement in Quantum Field Theory* (2nd ed.), Volume 5, pp. 335–345. Elsevier Academic Press Inc.

Fraser, D. 2008. The fate of 'particles' in quantum field theories with interactions. *Studies in History and Philosophy of Science Part B: Studies in History and*

Philosophy of Modern Physics 39(4): 841–859 .

Fraser, D. 2020. The development of renormalization group methods for particle physics: Formal analogies between classical statistical mechanics and quantum field theory. *Synthese*: 1–37 .

Fraser, D. 2021. Particles in quantum field theory, In *The Routledge companion to philosophy of physics*, eds. Knox, E. and A. Wilson, 323–336. Routledge.

Fraser, D. and A. Koberinski. 2016. The higgs mechanism and superconductivity: A case study of formal analogies. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 55: 72–91 .

Gale, E. and M. Zych. 2023. Relativistic unruh-dewitt detectors with quantized center of mass. *Phys. Rev. D* 107: 056023. https://doi.org/10.1103/PhysRevD.107.056023 .

Grimmer, D. 2023. The pragmatic qft measurement problem and the need for a heisenberg-like cut in qft. *Synthese 202*(4): 104 .

Grimmer, D., B.d.S.L. Torres, and E. Martín-Martínez. 2021. Measurements in qft: Weakly coupled local particle detectors and entanglement harvesting. *Phys. Rev. D* 104: 085014. https://doi.org/10.1103/PhysRevD.104.085014 .

Gross, D.J. and F. Wilczek. 1973. Asymptotically free gauge theories. i. *Physical Review D 8*(10): 3633 .

Grundling, H. and G. Rudolph. 2013. Qcd on an infinite lattice. *Communications in Mathematical Physics 318*(3): 717–766 .

Halvorson, H. 2001. Reeh-schlieder defeats newton-wigner: On alternative localization schemes in relativistic quantum field theory. *Philosophy of Science 68*(1): 111–133. https://doi.org/10.1086/392869 .

Halvorson, H. and R. Clifton. 2002. No place for particles in relativistic quantum theories? *Philosophy of Science 69*(1): 1–28. https://doi.org/10.1086/338939 .

Jäkel, C.D. and F. Robl. 2014. The relativistic kms condition for the thermal n-point functions of the $p(\phi)_2$ model.

Kapusta, J.I. and C. Gale. 2007. *Finite-temperature field theory: Principles and applications.* Cambridge university press.

Kijowski, J. and G. Rudolph. 2005. Charge superselection sectors for qcd on the lattice. *Journal of Mathematical Physics 46*(3). https://doi.org/10.1063/1.1851604 .

Koberinski, A. 2019. Parity violation in weak interactions: How experiment can shape a theoretical framework. *Studies in History and Philosophy of Modern Physics* 67:

64–77 .

Koberinski, A. 2021. Mathematical developments in the rise of yang–mills gauge theories. *Synthese 198*(Suppl 16): 3747–3777 .

Koberinski, A. 2023. What good is haag's no-go theorem? what axiomatic methods can teach us about particle physics.

Koberinski, A. 2024. Phase transitions and the birth of early universe particle physics. *Studies in History and Philosophy of Science* 105: 59–73. https://doi.org/https://doi.org/10.1016/j.shpsa.2024.03.006 .

Koberinski, A. and C. Smeenk. 2024, 01. Philosophical issues in early universe cosmology.

Landsman, N. 1988. Non-shell unstable particles in thermal field theory. *Annals of Physics 186*(1): 141–205 .

Le Bellac, M. 2000. *Thermal field theory.* Cambridge university press.

Malament, D.B. 1995. In defense of dogma: Why there cannot be a relativistic quantum mechanics of (localizable) particles. *University of Western Ontario Series in Philosophy of Science*: 1–10 .

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): 801–820 .

Miller, M.E. 2021. Worldly imprecision. *Philosophical Studies* 178: 2895–2911. https://doi.org/10.1007/s11098-020-01591-z .

Narnhofer, H., M. Requardt, and W. Thirring. 1983. Quasiparticles at finite temperatures. *Communications in Mathematical Physics* 92: 247–268 .

Papageorgiou, M. and D. Fraser. 2024. Eliminating the 'impossible': Recent progress on local measurement theory for quantum field theory. *Foundations of Physics 54*(3): 26 .

Peskin, M.E. and D.V. Schroeder. 1995. *An introduction to quantum field theory.* Advanced Book Program, Westview Press.

Politzer, H.D. 1974. Asymptotic freedom: An approach to strong interactions. *Physics Reports 14*(4): 129–180 .

Ruetsche, L. 2011. *Interpreting quantum theories.* Oxford University Press.

Ruetsche, L. 2024. The perils of approximate ontology. *Synthese 204*(4): 122. https://doi.org/10.1007/s11229-024-04777-8 .

Ruiz de Olano, P., J.D. Fraser, R. Gaudenzi, and A.S. Blum. 2022. Taking approximations seriously: The cases of the chew and nambu-jona-lasinio models. *Studies in History and Philosophy of Science* 93: 82–95. https://doi.org/https://doi.org/10.1016/j.shpsa.2022.02.013 .

Satz, H. 2011. The quark-gluon plasma – a short introduction. *Nuclear Physics A* 862-863: 4–12. https://doi.org/https://doi.org/10.1016/j.nuclphysa.2011.05.014 .

Sebens, C.T. 2022. The fundamentality of fields. *Synthese 200*(5): 380 .

Smeenk, C. 2019. Gaining access to the early universe, In *Why trust a theory*, eds. Dardashti, R., R. Dawid, and K.P.Y. Thébault, 315–335. Cambridge University Press.

Swanson, N. 2020. How to be a relativistic spacetime state realist. *The British Journal for the Philosophy of Science 71*(3): 933–957 .

Swanson, N. 2024. Three field ontologies for qft.

Teller, P. 1997. *An interpretive introduction to quantum field theory.* Princeton University Press.

Tumulka, R. 2022. *Foundations of quantum mechanics*, Volume 1003. Springer.

Wald, R.M. 1994. *Quantum field theory in curved spacetime and black hole thermodynamics*. University of Chicago press.

Wallace, D. 2001. Emergence of particles from bosonic quantum field theory. *arXiv preprint quant-ph/0112149* .

Wallace, D. 2020. On the plurality of quantum theories: Quantum theory as a framework, and its implications for the quantum measurement problem, In *Scientific Realism and the Quantum*, eds. French, S. and J. Saatsi, 78–102. Oxford University Press. https://doi.org/10.1093/oso/9780198814979.003.0005.

Zinn-Justin, J. 2000. Quantum field theory at finite temperature: An introduction.