



Particles, fields, and the ontology of the standard model

Federico Benítez¹ 

Received: 22 July 2022 / Accepted: 30 December 2022
© The Author(s) 2023

Abstract

In this work we discuss issues of ontological commitment towards one of the most important examples of contemporary fundamental science: the standard model of particle physics. We present a new form of selective structural realism, which uses as its basis the distinction between what have been called framework and interaction theories. This allows us to advance the ongoing debate about the ontological status of (quasi-)particles and quantum fields, by emphasising the distinction between quantum field theory serving as a framework, and the standard model itself, which we argue is an interaction theory embedded within this framework. Following a discussion of what ontological commitments corresponds to each of these two classes, we argue that some of the previous proposals in the literature might have been misguided by the blending of quantum field theory and the standard model into an undifferentiated unity, and defend a moderate form of object realism with respect to particle-like entities.

Keywords Scientific metaphysics · Standard model · Quantum field theory · Particle ontology · Field ontology

1 Introduction

The issue of ontological commitment is central to any kind of scientific realism, the contention that our best science refers to—has some kind of direct connection with—mind-independent reality. For the scientific realist, the ample success of science at predicting and explaining phenomena supports a metaphysics based on our best theories.

Nonetheless, not all theories deserve the same kind of ontological commitment, and a version of scientific realism that is truly consistent with science should also include information that can be read from the form or character of different theories. Namely, scientific theories come in different forms and shapes, and this is true even for

✉ Federico Benítez
federico.benitez@unibe.ch

¹ University of Bern, Bern, Switzerland

so-called fundamental theories—and this has to have consequences when assigning ontological commitments.

In this work we elaborate on these ideas, building up a proposal for ontological commitments in the case of what is arguably our best theory in fundamental physics: the standard model of particle physics (which we abbreviate by SM going forwards). As it is well-known, the predictions of this model have been tested to some of the highest levels of accuracy known to science today, and they have stood the test of almost half a century of experimental probing, culminating with the discovery of the Higgs particle at the LHC.

The SM represents one of the most important success stories for science. It is the result of a decades long collaboration of hundreds of theoreticians and experimenters (Kragh, 2002). Its development also marked a change in the way theoretical physicists approached their subject: compared to the generation of Einstein, Bohr, and Heisenberg, for example, younger physicists were much less interested in the philosophical implications of their discoveries, and paid much less attention to elaborated metaphysical or epistemological reasoning, beyond a somewhat naïve instrumentalism (Galison, 1997).

Traces of this ‘practical’ approach are easily detected within the theory, and complicate the work of building a consistent ontology. On top of that, there is the nagging issue of the measurement problem, which although seldom discussed in this context, quantum field theory inherits from quantum mechanics. Any proposal for an ontology of a quantum theory has to come to terms with the measurement problem, as measurements are where the link between the theory and the world is established. As is well known, there is little consensus regarding a solution to the measurement problem in QM. However, this should not hamper us to attempt a metaphysical examination of the SM. Even if we remain agnostic with respect to the measurement problem, there is enough information—both theoretical and empirical—to attempt to discern the ontology that suits the SM the best. This is what we do here.

The SM is an example of a quantum field theory (QFT from now on): a theory of fields which is compatible both with the precepts of quantum mechanics and with those of special relativity. But a clear distinction should be made between the framework of QFT and the specific interaction theory of the SM, set within such a framework. In particular, we want to emphasise that the constraints on the SM set by QFT and the predictions of the model itself do not share the same amount of empirical success, and generally speaking that the success of the SM does not necessarily translate to a blind confidence on the QFT framework. Indeed, the fact that an interaction theory is set within a framework is usually historically contingent, as can be seen from the example of Maxwell’s theory, initially thought to be contained within Newtonian mechanics, but ultimately shown to be naturally embedded within special relativity.

In order to decrypt what this means in a discussion of realism about modern particle physics, we have to first introduce our version of selective scientific realism, based on the difference between theories that serve as frameworks, as opposed to theories that describe specific interactions. We do this in Sect. 2. Then, in Sect. 3, we give some context on the ongoing discussion about the reality of particles and quantum fields. A central role is played by the use of renormalisation group techniques in QFTs, which we discuss in some detail in Sect. 4, as well as by gauge symmetries, discussed in

Sect. 5. With all this in place we give an account on our ideas on what being realist about the SM should look like in Sects. 6 and 7. We present our conclusions in Sect. 8.

2 Framework and interaction theories

We contend that a consistent selective scientific realism should take into account the ontological level of the difference between what have been called *interaction*, as opposed to *framework*, theories. This classification is a refinement proposed by Flores (1999) to the classical distinction made by Einstein between principle and constructive theories (Einstein, 1919).

Flores' classification revolves around a functional criterion: interaction theories are those that deal with the different ways in which entities are observed to interact (or somehow affect each other) in the world, whereas framework theories provide general constraints and common regularities for (generally more than one) interaction theories. A rich example of how this classification works can be studied at the level of Newtonian mechanics: Newton's three laws of motion give a regulative background for the study of any force law, and should be seen as the overarching or framing structure of classical mechanics. Conversely, specific force laws, such as the law of universal gravitation, deal with the details of one particular type of interaction, and work embedded within the general framework.

On the one hand, interaction theories directly deal with specific ways objects are seen to interact in the world. Examples of these include theories of the four fundamental forces, but also non-fundamental theories for phenomena such as elasticity, friction, or fluid mechanics. All these theories assume the existence of objects in the world, which participate in the interactions involved in each case. On the other hand, framework theories deal with general constraints to any such interaction, and they are developed in an indirect fashion, by examining the general structure of all interaction theories. Examples include Newton's three laws, thermodynamics, special relativity, and also, as we argue in more detail below, standard quantum mechanics.

The distinction between framework and interaction theories shares some characteristics, but transcends, the traditional distinction between kinematics and dynamics, or the difference between laws and meta-laws (see Lange, 2007a; 2007b). It differs from the kinematics/dynamics distinction in that framework theories, whereas including all the kinematics, often also deal with the general dynamic laws that all interactions have to obey, such as, e.g., Newton's second law, or Schrödinger equation. It differs from the law/meta-law classification in that frameworks not only deal with general symmetry principles but also with constraints to any interaction that do not seem to be a manifestation of a symmetry. An example of this kind of constraint is given by Born's rule in quantum mechanics, which we discuss further below.

The crux of the argument, as presented initially by Romero-Maltrana et al. (2018), and further elaborated by Benitez (2019) and Maltrana et al. (2022), is that there are distinctions to be made when assigning ontological commitments to the different objects, properties, and laws within each of these two classes of theories, a question that has not been properly discussed in the literature—and which in our opinion has

led in some occasions to confused metaphysical postures. The claim is that taking this classification into account has implications when reading the ontology of scientific theories.

Interaction theories by necessity assume the existence of interacting objects, the nature of which can only be accessed by means of the corresponding interaction in such a way that these objects are individuated by their position within the network of interactions in the world. Framework theories, on the other hand, are compatible with ‘empty worlds’ made purely out of the structure behind the constraints they encode, a structures-first view that would be natural for defenders of ontic structural realism.

All interactions involve some kind of ‘charge’, a property that serves as the source of the interaction. Our approach assigns ontological weight to the theoretical entities carrying these charges, even though admitting that we cannot have direct access to their intrinsic nature beyond this charge carrying capacity. Indeed, because of this inaccessibility, the most parsimonious metaphysical position to take with respect to these objects is to consider them simply as *bundles* of these interaction charges: nothing more than these properties instantiated together in space and time (see Paul, 2017, for a brief introduction to modern bundle theory). This parsimonious position will play a role in what follows.

Conversely, the objects (usually abstract objects) belonging to framework theories play the role of enforcing the general principles of the theory, a purely structural role. Framework theories provide constraints and spaces of possibilities for ‘empty worlds’, such as, e.g., Minkowskian space-time in special relativity, worlds which are only populated when one considers a specific interaction theory embedded within (e.g., electrodynamics). Framework theories do involve objects, such as e.g., space-time points or Hilbert space operators. But these are precisely the kinds of objects to which we struggle to commit ontologically, as their nature seems purely structural/mathematical. Framework theories thus accommodate the precepts of a structure-first ontology, and their objects are then best understood as ‘nomological’ entities, i.e., objects that are ontologically secondary to the structures they encode.

We now use this approach to study the ontology of the standard model of particle physics, which we contend is an interaction theory embedded within the framework of quantum field theory. As a first step we analyse this framework character.

3 Quantum field theory

3.1 The character of QFT

Quantum field theory is our best attempt for a relativistic version of quantum mechanics. It can be viewed as the outcome of imposing the constraints of quantum mechanics (QM)—such as Born’s rule and the Hilbert space description of physical states—together with those of special relativity (SR)—the structure of the Poincaré group—as well as extra conditions forcing strict space-time locality for observables. The notion of field seems to be essential to maintain the locality properties of SR, already at the classical level. However, it is important to notice that classical fields such as Maxwell’s electromagnetic field and the fields in QFT are very different kinds of objects, with

the latter being *operator valued* (more on this below), which makes a straightforward ontological commitment to these extremely abstract objects at the very least doubtful.

The combination of the constraints of SR and QM severely reduces the space of possible physical models for relativistic quantum systems, and QFT serves as a framework for these. Explicitly, we contend that, because SR and QM are both framework theories, QFT must also be a framework. The case of SR is quite clear: the theory starts from two general postulates assumed valid for any physical process, serving as a framework for interaction theories such as electromagnetism. In fact relativity was the theory motivating Einstein to introduce his distinction between principle and constructive theories, which serves as a basis for the framework/interaction distinction. As mentioned above, the world of SR is ‘empty’ until objects belonging to some interaction theory—originally electrodynamics—are brought in. What SR does is to put constraints on any and all interacting matter existing within this world.

In fact, the same concepts can be used to describe how QM works as a theory. It also starts from a series of postulates about how physics works for any system, which frame theories about e.g., atomic structure or condensed matter. Examples of these postulates include the Born rule (and more generally the fact that the outcomes of measurements are in general probabilistic), and the Hilbert space structure behind the description of quantum states. Within this general framework, which constraints the behaviour of any physical system, one uses interaction theories to describe specific systems, such as, again, electrodynamics to describe atomic structure. The top-down framework structure of QM can most easily be appreciated in the diverse approaches aiming for an axiomatic (Moretti, 2018) or informational theoretical (Clifton et al., 2003) expression of the theory.

As a theory that is mostly used to build models of microscopic phenomena, it is not easy to disentangle QM from its applications in atomic and molecular physics, but the fact is that QM gives a framework that is expected to be valid at any scale and for any physical system (e.g., including macroscopic measurement devices). QM constrain how observations work for any physics involved, be it atomic, condensed matter, or eventually also at the level of the whole universe, such as in the field of quantum cosmology.

The framework character of QM is independent of which interpretation of the theory is chosen to deal with the measurement problem. It is true that some interpretations of QM seem to include a type of interaction used to explain quantum behaviour, but we argue that these purported interactions are not really part of QM. For example, the standard Copenhagen interpretation includes wave function collapse upon measurement of a quantum system. This collapse may be considered as a mark of some interaction taking place, most reasonably some interaction between the system and an experimental device. This interaction, however, is ultimately described by an independent theory, an interaction theory, such as the electromagnetic interaction for the purposes of measuring an electron’s position. The fact that within this interpretation we don’t get a detailed description of how the electron wave function collapses when interacting with the device does not mean that a new kind of interaction is needed.

Following this short discussion about the character of QM and SR, it is easy to deduce that if QM and SR are framework theories, so too should QFT, which basically is the fruit of their combination. This framework theory is to be distinguished from

the SM, i.e., the interaction theory of three of the four known fundamental forces. Of course, the SM has to be understood within the framework of QFT, so it is difficult to disentangle the two theories in order to find a consistent ontology, and in particular an ontology of subatomic entities. We propose a possible solution in what follows, but first let us discuss the present state of the literature on the subject.

3.2 QFT and particles

There is an ongoing debate in the literature about the correct ontology for QFT, the question being if an ontology in terms of fields should be preferred over an ontology in terms of particles (for an overview of the main disagreements see Egg et al., 2017). The issue arises mainly because the version of QFT best suited for formal mathematical analysis, algebraic QFT (Araki, 1999; Halvorson & Müger, 2006; Haag, 2012) provides a number of rigorous results showing that the notion of localised particle is impossible to represent within an interacting algebraic QFT. The most famous of these results is Haag's theorem, which states that, as opposed to a free theory, an interaction theory does not have a Fock space representation, that is, a representation in terms of particle number. In intuitive terms these results show that even though it is possible to have a notion of a localised electron in a theory of electrodynamics with the interaction coupling set to zero (a free theory), as soon as an interaction is present this notion is no longer available.

This state of affairs is in stark contrast to non-relativistic quantum mechanics, where there is always a unitary equivalence between the representation of particles in a free theory and that on an interacting theory. This fact is used to construct what is known as the interaction picture, where operators are evolved using a free field representation while states evolve using the interacting field representation—with many applications in perturbative approaches. Within the formalism of algebraic QFT such a picture formally does not exist, because there is no such an equivalence. Importantly, as we discuss in more detail below, the existence of inequivalent representations is an issue that not only affects the chances of finding an ontology of particles for QFT, but also an ontology in terms of fields (Baker, 2009).

Another well-known challenge to the particle picture within QFT is the Unruh effect (for recent discussions see Ruetsche, 2002; Earman, 2011; Ruetsche, 2012). This effect is arguably not embedded within the proper QFT framework, as it takes place when one compares the vacuum state a field theory (the ground states with no particles) as observed by two different observers in accelerated motion with respect to each other. In this way, at least one of the observers is not an inertial observer. What one can show is that the vacua for both observers are inequivalent, so that in some cases the vacuum of one observer is not even in the space of quantum states of the other. This comes about because—again—the two vacua lead to unitarily inequivalent representations.

The shocking aspect of the Unruh effect is that from the viewpoint of the accelerating observer, the vacuum of the inertial observer will look like a state containing many particles in thermal equilibrium—a warm gas of particles whereas the inertial observer measures nothing. This very counter-intuitive effect seems to indicate that the defini-

tion of what constitutes a ‘particle’ depends on the state of motion of the observer. As an aside, it is interesting to notice that there is a curved space-time analogous to the Unruh effect (as one would expect from general relativity), which is the phenomenon of Hawking radiation (Carroll & Ostlie, 2017): thermal radiation originating near a black hole event horizon, associated with the creation of particles in the region.¹

As a final example of the issues one finds when defining relativistic theories of quantum particles we briefly discuss Malament’s theorem (Malament, 1996; Halvorson & Clifton, 2002). Malament shows the inconsistency of a few intuitive desiderata for a relativistic, quantum-mechanical theory of particle mechanics. The treatment here is quite general, and shows that a relativistic quantum theory cannot define a completely localised observable, such as would be the presence of a particle on a given point. Notice however that there are loopholes (see for example Barrett, 2002; Oldofredi, 2018), as is the case with any theorem. In particular, at the most trivial level, non-relativistic quantum mechanics can also lead to results like these, as the quantum wave function cannot stay perfectly localised, even when we deal with the wave function for a point particle. This does not necessarily mean that particles do not belong to the ontology, only that the framework constraints how a particle can be represented.

All these examples show the conceptual difficulties associated with embedding the notion of particle within a theory that has to follow all the constraints of both special relativity and quantum mechanics. These are serious issues, even if their existence has not prevented physicist to very successfully model experiments performed in particle accelerators, where the most relevant processes are the collision and the decay of subatomic particles. Of course, there are many levels of the discussion, and the one we are interested in, the ontological level, is not really about the empirical success of a concept such as the particle concept. But it is also not, or not only, about abstract structures, such as the C^* -algebras that we use to encode the constraints we observe in nature. This difficult-to-transit ontological middle ground is the one that interest us the most here.

Nonetheless, these issues have led some prominent philosophers to deny the possibility of an ontology of particles for the SM (see e.g., Earman and Fraser, 2006; Fraser, 2008; 2009; 2011; Bigaj, 2018). The usual argument goes as follows: given that QFT is part of our best current physics, it should be the base of our ontology. As QFT does not allow the existence of particles in interacting theories, then we can conclude no such entities exist in the world. Thus, it is argued, an ontology in terms of fields is to be preferred. In turn, this has been used as an argument for the elimination of objects from the ontology of the world, as argued by the proponents of eliminativist ontic structural realism: as the argument goes, the fact that we cannot define what a particle is in our best theory of particle physics would be a strong indication that particles are not a real object in the world (Ladyman, 1998; Ladyman & Ross, 2007; French, 2014; Glick, 2016; Berghofer, 2018).

There are several problems with this view. While it is true that algebraic QFT does not allow for a particle representation, the ‘fields’ of QFT are quantum operator valued,

¹ For more technical details about Haag’s theorem, inequivalent representation, the Unruh effect, and related issues, we refer the reader to (Ruetsche, 2002; Earman, 2011; Ruetsche, 2012; Egg, 2014; Kuhlmann & Stöckler, 2018).

and they belong to an abstract infinite dimensional Hilbert space, not to physical space-time. The entities that would reasonably play the role of an ontology in terms of fields would be sums over possible field configurations—a quantum superposition of classical fields. These objects are also problematic at the formal level, as they run into most of the same issues affecting the particle interpretation, including the issue of inequivalent representations. In fact, the Fock-space representation and the sum over field configuration representation are strictly equivalent (see Baker, 2009). Also, algebraic QFT not only does not allow the existence of local particle-like states, but neither of local observables, that is to say, locally defined measurement outcomes, which is problematic (to say the least) for any theory that pretends to have some contact with empirical observations (Halvorson & Müger, 2006). In addition to this, and for reasons to be discussed below, the natural candidates for ontological states are gauge invariant quantities, an invariance which is not generally complied by quantum fields configurations.

Besides this, even though it is the most formally coherent version of QFT, algebraic QFT still cannot deal—barring low dimensional models (see Summers, 2012)—with interacting theories, nor with gauge symmetries (at least not with the fundamental step of gauge fixing *à la* BRST (Weinberg, 1996)), nor, most importantly, with the renormalisation group. The renormalisation group is fundamental to understand e.g., running couplings, which lead to asymptotic freedom in the case of quantum chromodynamics (QCD), a most relevant part of the SM. It also plays a fundamental role in the search for alternatives and extensions to the SM. One could go as far as to say that to study particle physics and the SM without the tools of the renormalisation group is beyond hopeless. If the scientific realist wants to use our best physics to construct a metaphysics, she should for now turn to the ‘conventional’ Lagrangian version of QFT, with all the formal pitfalls that this might entail. Notably, this version of QFT is regularly used to model particle phenomena, albeit in most cases by using some sort of approximation technique.

Lagrangian QFT is a very loose term. Here we take it to mean the path integral representation of a QFT (Weinberg, 1995). Explicitly, given a Lagrangian density $\mathcal{L}[\Psi(x)]$ as a functional of a set of space-time fields collectively denoted by $\Psi(x)$, one defines the partition functional as a weighted sum over arbitrary fields configurations: a path integral. This functional is usually defined as the continuum limit of a lattice field theory (Weinberg, 1996), although, we insist, it is formally well-defined only in some special cases (such as for a free theory). From the partition functional one can deduce all the correlation functions of the theory. The partition functional encodes in this way all the information about the theory. In applications in particle physics, for example, the 4-point correlation function can be associated with collision processes with two incoming and two outgoing particles.

Quantities associated with particle states can be formally defined for a Lagrangian QFT, but their computation is usually far from trivial. Many techniques have been developed to approximate them, the most well-known being the perturbative expansion in terms of Feynman diagrams. In this regard, it is important to emphasise that Lagrangian QFT goes much beyond a technique for obtaining collision cross sections or particle lifetimes, and in particular to differentiate Lagrangian QFT from the perturbative approach as presented in particle physics textbooks. The ontologically

relevant aspects of Lagrangian QFT are to be searched within the exact results that can be obtained with the greatest generality, and not in the technically admirable but approximate and much less formally satisfactory results stemming from Feynman diagrammatics.

The framework character of Lagrangian QFT allows for the demonstration of a number of very important constraints to be fulfilled by any possible theory of particle physics (at least as understood up to now). These constraints include the necessary existence of an anti-particle for each type of particle, the spin-statistics theorem, the CPT invariance theorem, and the existence of a Källen-Léhman representation for the particle propagator, of which more is discussed below. But by far the most relevant aspect of Lagrangian QFT for the deduction of general constraints is that it allows to employ the machinery of the renormalisation group. Before seeing delving deeper into questions of ontology, let us discuss some important, and not much explored aspects related to this powerful technique.

4 Lessons from the renormalisation group

Since the works of Wilson in the seventies (Wilson & Fisher, 1972; Wilson & Kogut, 1974; Wilson, 1975) it is impossible to disentangle the properties of Lagrangian QFT from those of the renormalisation group (RG). One can make a good case that the development of RG techniques has been the last big breakthrough in theoretical physics, leading to a unified description of a wide range of a priori unrelated phenomena—a hallmark sign of a framework. In this section we discuss some of the ontological implications one can take from the RG framework. For a recent similar take see Williams (2019).

Even so, reliance on the RG has been posed as one of the disadvantages of the theory in the literature: Fraser (2009) explicitly states that the approaches based on the renormalisation group depend on regularisation procedures that break Lorentz invariance (and therefore special relativity), and are thus incompatible with the tenets of a full theory of QFT, which would be best approached from first principles in the algebraic spirit. But this is simply not true beyond the most basic textbook examples: it is always possible (though perhaps technically highly non trivial) to find regularisation schemes that do not break any of the symmetries of the system (Berges et al., 2002; Delamotte, 2012). Even in the usual textbook perturbative treatment, it can be argued that Pauli-Villars regularisation or dimensional regularisation (the most widely used in the literature) do not break Lorentz invariance (Weinberg, 1996).

It is nonetheless the case that within the philosophy of physics, ‘renormalisation’ has some bad connotations, being associated with a set of informal semi-phenomenological manipulations developed in order to deal with divergences in field theories. This is very much at the origin of the method, for example in Bethe’s computation of the Lamb shift (Kragh, 2002). However, the modern theory of the RG does not deal with infinities, or even necessarily with field theories: it is instead associated to studying how theories change under scale transformations (Delamotte, 2002, 2012; Butterfield & Bouatta, 2015), and its relevance goes much beyond particle physics—with statistical mechanics and condensed matter physics being additional examples of

its wide range of application. As Delamotte (2002) discusses while reflecting about the history of the field

...the divergences of perturbation theory in QFT are directly linked to its short distance structure which is highly non-trivial because its description involves the infinity of multi-particle states (...) However, strangely (at least at first sight) the theoretical breakthrough in the understanding of renormalisation beyond its algorithmic aspect came from Wilson's work on continuous phase transitions. The phenomena that take place at these transitions are neither quantum mechanical nor relativistic and are nontrivial because of their cooperative behaviour, that is, their properties at large distances.

In the case of particle physics, a Lagrangian QFT is defined by its field content, the symmetries of the theory, and a (possibly infinite) set of couplings related to physical parameters such as particle masses and interaction strengths. The values of these couplings have to be measured experimentally, and are defined at an energy scale μ associated to the scale where the measurements of the coupling takes place. In the modern framework, the *renormalisability* condition simply states that knowing the couplings at a certain energy/length scale μ is sufficient to characterise the theory at any other, arbitrary, energy scale. The RG provides a precise description of how the couplings change when they are observed at different scales. As an example, the fine-structure constant in quantum electrodynamics, $\alpha = \frac{1}{137}$ at zero energy, but when measured at energies corresponding to those in particle accelerators the 'constant' is seen to be greater, e.g., $\alpha \simeq \frac{1}{125}$ at an energy level of 10^{12} eV, a value which is correctly predicted by the RG.

Perhaps the most important lesson that should be taken from the RG is what is known as the decoupling of high energy (equivalently, short distance) modes. Plainly speaking, this property means that QFTs are in general insensitive to the detailed physics at high energies/short distances. The reason for this is that the general RG equations are local in the energy scale μ , and therefore tend to isolate different energy scales. On top of that, when one solves the RG equations from higher to lower energies, couplings related to interactions between a large number of fields, which can become important at higher energies, usually go to zero for dimensional reasons. This decoupling property is well known in the study of critical phenomena in statistical mechanics, where one knows that the behaviour of the system at large distances does not depend on the details at the level of the lattice (e.g., the system behaves as isotropic even though it is defined in a lattice which is not symmetric with respect to rotations, and it is not relevant for the long distance properties if the underlying lattice is triangular or rectangular, etc.).

The quantum theory of electrodynamics, QED, brings about a great example of this decoupling property. In its simplest form, the theory depends only on a matter field (say for the electrons/positrons), together with a dynamical $U(1)$ gauge field associated with the photons. This is a renormalisable theory, well-defined up to extremely high energy scales. This means that QED is completely insensitive to the physics happening at higher energies, and in particular, the theory cannot tell us anything about what happens at the (relatively low) electroweak energy scale, where the Weinberg-Salam-Glashow electroweak model gives the correct description of the observed physics, and

QED stops being a good representation of reality. The physics at these scales include very important effects such as spontaneous symmetry breaking and the generation of all fundamental masses via the Higgs mechanism, and QED is completely oblivious to all this. In a way, renormalisability is both a blessing and a curse, as one could be tricked to believe that the theory is valid up to much higher energy scales than it really is.

That is to say: the theory of QED is valid at energies much higher than those of its physical relevance. Mirroring the pessimistic meta-induction argument, it is very likely that the SM, also a renormalisable theory, is oblivious to new, as-yet unknown physical effects at energy scales higher than our current experimental limits. From this point of view, all QFTs should be seen as *effective* theories only, valid up to a certain energy scale. This range of validity is much easier to see in what are traditionally known as non-renormalisable theories, such as the Fermi theory for weak interactions: the theory breaks down at an energy scale of the order of magnitude of the electroweak symmetry breaking, which is the scale where its effective description of electroweak processes such as beta decays ceases to be valid, and one has to use the full Weinberg-Salam-Glashow Model instead.

This is not only very well-known to physicists; it is in fact behind one of the most common approaches to explore possible physics beyond the SM, what are known as effective field theories (Georgi, 1993; Costello, 2011). Broadly speaking, these work as follows: a number of (technically non-renormalisable) extra terms are added to the Lagrangian of the SM to model the most probable (the most ‘relevant’ in the RG sense) corrections to the SM, leading to a number of predictions for high energy physics differing from the ones in the SM in ways that are amenable to be tested experimentally.

Thus, even though historically the renormalisation procedure was introduced to deal with infinities appearing in QFT calculations, the modern theory of the RG is a different beast altogether, as stated by Delamotte (2002) in his pedagogical overview:

...although the renormalisation procedure has not evolved much these last thirty years, our interpretation of renormalisation has drastically changed: the renormalised theory was assumed to be fundamental, while it is now believed to be only an effective one; Λ was interpreted as an artificial parameter that was only useful in intermediate calculations, while we now believe that it corresponds to a fundamental scale where new physics occurs; non-renormalisable couplings were thought to be forbidden, while they are now interpreted as the remnants of interaction terms in a more fundamental theory. Renormalisation group is now seen as an efficient tool to build effective low energy theories when large fluctuations occur between two very different scales that change qualitatively and quantitatively the physics.

Remarkably, an effective field theory can give a reliable description of a physical system even if the underlying physical degrees of freedom are not fields. This is due to the property of decoupling of short distance modes that we described above. As a classic example there is the famed XY model (Le Bellac, 1991; Cardy, 1996), which

consist of a 2 dimensional lattice of 2 dimensional vector spins with a tendency for alignment, and a Lagrangian² that is symmetric with respect to global spin rotations.

This system shows a transition (the Berezinskii-Kosterlitz-Thouless (BKT) phase transition) from a high-temperature disordered phase to a low-temperature quasi-ordered phase. A closer study shows that this phase transition is controlled by the dynamics of vortices, as defined by the circulation of the spin vectors in the lattice. In the XY model vortices are topologically stable configurations of spins. It is found that the low-temperature ordered phase vortices bound together forming vortex-antivortex pairs, whereas in the high-temperature disordered phase vortices are free.

A RG study of the BKT transition deals with the behaviour of an effective field theory with the same symmetry group as the XY model, so that the fundamental spin vector degrees of freedom are completely ignored. This means that the field theory cannot take into account a dynamics of vortices made from the spin vector spins. Remarkably then, even without taking the vortex degrees of freedom into account, the RG method can explain the characteristics of BKT order (for a modern example see Jakubczyk et al., 2014). This is a clear example of how, even though the exact fundamental degrees of freedom are ignored—either because their inclusion is technically challenging or because we do not have access to them—a field theory can yield the relevant low energy information about a system.

Indeed, one can say with a good degree of confidence that the QFT picture must break down at some point—at the very least at the Plank scale, as taking SR at face value at these energy levels would represent an excess of optimism. The QFT framework is only validated by the low energy information available from the behaviour of interaction theories, and we lack this information for ultra high energies.

In summary, Lagrangian QFT, as a framework theory, should properly be seen as giving a blueprint for constructing families of effective theories, as described by the RG, and not as an ultimate map of reality.³ These effective theories are by construction not expected to correctly describe the fundamental level of reality. Ontological commitment to the entities belonging to such a framework for effective theories is thus unjustified from several fronts. The issue stands however that the SM, one of our best theories up to now, stands within the framework of Lagrangian QFT. It is natural to demand if it is then even possible to deduce a coherent ontology of the world from this piece of fundamental science. In what follows we try to give a tentative answer to this.

5 Gauge theories and realism

The SM is a very specific type of Lagrangian QFT, namely a gauge theory—being invariant with respect to what is known as a gauge symmetry. Gauge symmetries refer to symmetries in some internal space, associated with a given symmetry group, specifically $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ for the SM. Generally speaking, symmetries constitute a powerful organising principle in science, and are a prime example of

² Technically, a Hamiltonian. But it plays the same role as the Lagrangian in QFT.

³ Here we oppose works like Kuhlmann (2013).

structures in our theories.⁴ Indeed, ontic structural realism proponents such as French (2014) claim that our ontological commitment should be put on structures such as symmetries and laws, and not on objects such as fundamental particles.

When analysing physics, it is useful to classify symmetries in two types: empirical and theoretical symmetries. This classification is justified as follows. A set of physical situations can possess a symmetry that can be detected empirically, whether or not we have a precise theoretical description of what is at work. One can thus see Galilean invariance as an empirical symmetry, as shown e.g., by the experience of a train moving at constant speed: physical processes taking place inside the train (or, as Galileo himself proposed, a ship) are indistinguishable from their analogous happening at rest with respect to the surface of the Earth.

The case of theoretical symmetries is more subtle. A theoretical description of a physical situation occurs by means of a mathematical *model* that represents the situation within a given theory. Under certain circumstances, it can be the case that more than one model of a theory describes the same physical situation. One says that a function between different models of a theory is a theoretical symmetry if it leaves invariant the underlying physical situation, that is, if it represents the same situation in the world. Changes of coordinates are the example *par excellence* of this kind of symmetry. Another (related) example is diffeomorphism invariance in general relativity.

This classification is important when considering the supposed modal role played by symmetries, as discussed within ontic structural realism, and also when dealing with the mathematical/physical distinction, which is vital in order to avoid bad metaphysics. Mathematically, symmetries are described by a particular kind of structure called a group, but we only have indirect access to these structures by means of the observed behaviour of interactions in the world.

Gauge symmetries, in their most general form, are field transformations that leave the Lagrangian of the theory invariant, and that are not space-time transformations (such as would be a rotation in space). Instead, they are transformations acting on intrinsic properties of the fields themselves. Here it is important to distinguish between global and local gauge symmetries. A global symmetry is one that acts the same way over all space-time, such as would be multiplication of the fields by a constant phase factor $e^{i\alpha}$. A local symmetry is one that acts independently on each space-time point, e.g., multiplication by a position-dependant phase such as $e^{i\alpha(x)}$.

Gauge symmetries clearly seem to follow the mold of theoretical symmetries (Healey, 2007)—but the situation is subtler than that, and a further distinction is needed. Following a theorem by Weyl (Weinberg, 1995), the energy spectrum (i.e., the masses of the particles) of a quantum theory represents the structure of the global symmetry group of the theory, i.e., when a Lagrangian is globally symmetric under a certain symmetry group this manifests in the spectrum of particles: as multiplets of particles with the same mass. This constitutes an empirical effect of global symmetries. Indeed, historically this is how the $SU(3)$ group was proposed to underlie quantum chromodynamics.

⁴ See Brading and Castellani (2003) for a detailed analysis of the role of symmetries in physics.

We contend that local gauge symmetries, on the other hand, do not have such empirical effects and are purely theoretical. This takes a lot of weight out of the purported modal role played by gauge symmetries. In this view, local gauge symmetries are not structures ‘out there in the world’ but merely theoretical constructions that allow our QFT mathematical models of particle physics to be simpler.

Lyre (2004), as well as Healey (2007) give a compelling analysis of the non-physicality of what is known as the gauge principle—that is to say, the ‘upgrade’ of global symmetries of a theory into local, space-time dependent, symmetries. Indeed, there is nothing forceful or natural about these local symmetries, which are by definition unobservable. This, of course, does not take away from their celebrated usefulness as theoretical tools. See also Belot (2003).

These issues cannot be solved by the mere addition of technical sophistication, as is sometimes attempted. For example, gauge field theories can alternatively be described by means of a mathematical technology called a fiber bundle. Basically, a fiber bundle assigns to any point of space time (the *base manifold*) what is called a *fiber*, a space isomorphic to the gauge group of interest. In this formalism, the gauge field A is identified with the geometrical *connection*, which allows to parallel transport through different fibers. Thus, the state of the system can be represented by a slice through this abstract space, and a gauge transformation can be seen as a change or deformation of the slice.

Even if one argues then that the structure of the world is best described as a fiber bundle with Minkowski space as a base manifold, and the fibers being isomorphic to the gauge symmetry group, a local, space-time dependent, gauge transformation is not physical in any way, but merely something resembling a change of coordinates in this abstract space (Healey, 2007).

Analogously, there is a sophisticated category theory analysis by Weatherall (2016) which argues that even if electromagnetism is a gauge theory which has excess mathematical structure when modelling the world, more general fiber bundle (Yang-Mills) theories and general relativity do not suffer from this fact. Here what is at stake is if the technical definition of excess structure as given by (Weatherall, 2015, 2016) is the physically relevant one. The fact remains that there is no known way to physically implement a gauge transformation, which seems thus to be a purely theoretical symmetry for all these models.

One can analyse in a similar fashion all the purported empirical evidence for local gauge symmetries. Be it ghost fields when performing gauge fixing (which are by definition purely mathematical convenience), gauge fixing within the Higgs mechanism, which again is merely a computational tool, or any other of the technical QFT arguments for the empirical import of local gauge fields; in the end one can always reach a conclusion that was already true by the time of Maxwell’s theory: gauge fields involve extra degrees of freedom and arbitrary choices that are unobservable and non-physical.

Perhaps the most compelling argument for the interpretation of gauge transformations as physical comes from the famed Aharonov-Bohm (AB) effect (Aharonov & Bohm, 1959). This is a purely quantum mechanical effect, in which, due to the direct coupling between the electromagnetic potential A_μ —a gauge field—and the quantum wave function ψ , detectable electromagnetic effects on the phase of the wave function can be observed, even in physical configurations, such as on the exterior of a solenoid,

where the magnetic field is vanishing. This phase shift can be detected as a movement of the interference pattern of a two-slit experiment.

We do not consider the AB effect as a strong argument for an ontology of gauge-dependent quantities. There are many possible alternative explanations for the effect that do not imply the reality of gauge fields (see, e.g., Vaidman, 2012; Lazarovici, 2018). The detailed discussion of these issues goes beyond the scope of this work, but suffice it to say that one can devise ways to avoid assigning reality to gauge fields, without affecting the observed predictions of AB.⁵ The moral would again be that a mathematical object that is prone to arbitrary modifications should confront a high hurdle for ontological commitment.

What all this points to is that the structure of gauge symmetries is basically a convenient way to describe the actual physical systems, and has no physical meaning by itself. This goes against the claims that the primitive components of the world are these symmetries.

As an additional, independent argument against a field ontology for the SM consider now the following. The presence of a gauge symmetry (or of any other symmetry) implies that there are certain transformations that can act on the underlying fields, but are not amenable to experimental detection, even in principle. This has important consequences when trying to build a field ontology.

The basic argument is an extension to theoretical symmetries of an argument that goes back to Lange (2001), where it is argued that Lorentz invariance should be a requirement for ontological commitment in relativistic theories. The reasoning behind this is that what is real should not depend on arbitrary choices made by the observer. In this spirit, it seems reasonable to demand real entities to be gauge invariant, as the extra degrees of freedom contained in gauge dependent quantities are (i) superfluous, and (ii) underdetermined by observations, even in principle. Within the theory, any observable must be gauge-independent, but the fields themselves are not. Insofar as fields are generically gauge-dependent, a realist position about quantum fields entails an unnecessary underdetermination of fundamental entities. Note that such underdetermination would be radical, being as it is a requisite for the correct functioning of the theory.⁶

The situation is different to that of classical fields, as the quantum theory is only defined in terms of the potential (gauge dependent) field, whereas classical electromagnetism can in principle be defined in terms of the gauge invariant electric and magnetic fields. The desired gauge independence can in principle be achieved within the language of quantum fields, by means of a gauge fixing procedure (such as the Fadeev-Popov method (Weinberg, 1996)). However, the gauge condition that has to be chosen to do so has itself a large degree of arbitrariness, and the gauge fixing procedure enters into the serious issue of the Gribov copies of a gauge orbit (Weinberg, 1996).

⁵ Vaidman argues that the semi-classical perspective that underlies standard presentations of the AB effect constitutes an erroneous starting point, and that a consideration of the quantum character of the whole experimental device dissolves the effect. Even more radically, Lazarovici argues against assigning ontological weight to *any* field, not only the vector gauge field involved in the AB effect. The primitive ontology would be one of particles interacting directly, by means of generalizations of the Wheeler-Feynman version of electromagnetism.

⁶ See also Redhead (2002) on the interpretation of gauge symmetries.

Thus, from a philosophical—as opposed to pragmatic—viewpoint, it would be more natural to use genuinely gauge invariant states for our ontology, as e.g., holonomies (as proposed by Healey, 2007), or, even more justified by our interaction theories, localised gauge-invariant states representing particles, or close approximations to particles.

We can now lay out all of the reasons why an ontology of the SM based on quantum fields does not work. First, quantum fields are very abstract mathematical objects, that assign quantum operators belonging to an infinite dimensional Hilbert space to each space-time point. If one is realist about these fields, it is highly unclear what are the physical objects in space-time that build up the ontology. The best option seems to be a quantum superposition of classical fields, but this take on an ontology of fields enters into the same issues as the Fock-space particle representation (Baker, 2009).

Secondly, the RG procedure clearly shows that QFT gives an *effective* description of any physical system at energy scales compatible with the rules of QM and SR, *independently* of what are the true underlying degrees of freedom. Indeed, we know from all the examples in condensed matter physics that the field description is not necessarily the fundamental one, which should give us pause also in the case of particle physics.

Third, in the SM at least, the quantum fields that enter the description of the model are gauge dependent. The problem is, gauge dependent entities should not form part of our ontology for the reasons discussed in this section. The ontology of the world should be based on gauge-invariant objects.

Most importantly, and to stress the main point of this work: Lagrangian QFT—in fact every version of QFT—is a framework theory. To claim that the fundamental degrees of freedom in nature are quantum fields equates to giving full ontological weight to objects belonging to a framework. This is, we claim, misguided. One should study the interaction theories framed inside QFT in order to assign ontological commitments. We do that in what follows.

6 Towards an ontology of the standard model

The semi-empirical construction of the SM clearly implies its character as an interaction theory. In fact, it is the (more or less comfortable) combination of *two* such theories: a theory of the strong nuclear force (QCD), together with a theory of the electroweak interaction, and its spontaneous symmetry breaking (the Weinberg-Salam-Glashow model). The theory of the SM stands within the framework of Lagrangian QFT. As such, as per the discussion above, it should be considered an example of an *effective* field theory, that is, a theory that approximates the physics within a certain range of energy scales, having no access to higher energy phenomena due to the generic decoupling mechanism described by the RG formalism.

What we contend here is that the SM is, at its core, a theory about particle interactions and decays. The fact that this theory has to fit inside the framework of both special relativity and quantum mechanics implies that it has to be described as an effective QFT. That is to say, when constructing the interaction theory of the SM one has to respect the constraints posed by the framework of Lagrangian QFT, independently of the actual characteristics of the underlying ontology. This would be analogous to the

case in condensed matter physics, where one can use field theory and renormalisation group techniques to model these systems, even though at the fundamental level they are made up not of fields, but of discrete microscopic elements such as electrons or spin degrees of freedom.

As is clear from the history of electrodynamics, the framework theory within which an interaction theory stands can turn out not to be well-suited to account for all the properties or laws of the interaction. In that case, the framework of Newtonian mechanics was not compatible with some of the properties of the interaction theory of electromagnetism, and a new framework—special relativity—was necessary in order to have a fully consistent description of electromagnetic phenomena. This can perfectly well also be the case for a theory of interacting particles standing within a QFT framework.

The fact that the SM sits within such a framework does not mean that this is the right one. As with Maxwell's theory within the framework of classical mechanics, an incorrect framework theory could lead to paradoxes, like in this case the impossibility of satisfactorily representing what empirically are described as particles undergoing particle interactions. As QFT is the combination of the constraints of QM and SR, this inadequacy could be due to the fact that either QM or SR are incomplete descriptions (which is known to be the case at least for SR, given the existence of general relativity, and arguably also for QM as per the measurement problem). The possible solutions could be manifold: a discrete space-time, a generalisation of Bohmian mechanics, or a full M-theory could provide the correct notion of particle, among other examples. Again, all these possibilities would behave like a QFT at the energy scales associated with SR and QM, which are the ones to which we have access up to now.

It is important to emphasise this point: any interaction theory of particles that is compatible with QM and SR can *necessarily* be written in terms of a Lagrangian QFT. For example, to make a theory of electronic interactions one needs to know that electrons are fermions, and then to measure (by means of a scattering experiment) what are the relevant parameters of the Lagrangian density, i.e., the electron mass and the fine-structure constant, couplings that are defined at the energy scale at which the scattering experiments take place. This does not imply at all that electrons should be given an ontology in terms of fields. This is because the QFT description, through necessary at these energy levels, is just an effective description, as the RG procedure shows. That means that one ignores the truly fundamental degrees of freedom *by design*. This is indeed a good feature of QFT, because scientists don't have empirical access to the energy scales associated with 'truly fundamental' physics, such as e.g., those at the Planck scale. No matter what the world looks like at these high energy scales, be it fields, particles, super strings, or something still to be imagined, it will always look like a QFT at the scales we have access to in current particle accelerators.

It is in this sense that one can interpret the words of Weinberg (1995)

The reason that our field theories work so well is not that they are fundamental truths, but that any relativistic quantum theory will look as a field theory at sufficiently low energy [that is, at the energy levels attained in current accelerators].

The phenomenology of the SM is clearly based on the notion of interacting *particles*, that is, discrete quanta of matter that undergo interactions, as well as creation and decay processes. These particles are trivially gauge-independent, they are localised in space-

time, and they can be counted. In order to cope with the limitations of the Lagrangian QFT framework, SM computations of such processes have to be performed by means of starting and ending ‘asymptotic states’: initial and final states that are so far away from everything else in the world that they can be considered as free particles, and therefore be given a proper particle interpretation within the formalism of QFT. That is to say, the SM as a theory of particles has to work around the limitations of the QFT framework, so that the connection of the theory with observations is made via many of these workaround strategies.⁷

In fact, the formalism of Lagrangian QFT has to be bent and deformed to extract from it the predictions of the SM. A lot of the most technically challenging aspects of the theory are precisely those dealing with how to model particle collisions and decays. This is what lays behind gauge fixing, the perturbative Feynman diagram technique, and the traditional and the modern renormalisation group methods, among other examples.

As an additional argument, we can also experimentally probe the existence of point particles independently of the SM. Indeed, observation of a single electron in a Penning trap suggests the upper limit of the particle’s radius to be 10^{-22} meters (Dehmelt, 1988). This means that an electron is indistinguishable from a point particle up to distances corresponding to an energy scale three to four orders of magnitude larger than the one at which the SM is tested in the LHC⁸, so physicist have higher quality evidence for the particle-like nature of electrons than for the rest of the SM. Penning traps, a type of electromagnetic trap for charged particles, also have been used to trap single protons and measuring their properties, and they behave very much like localised particles (Mooser et al., 2014). These measurements are made using some of the interactions in which electrons and protons take part, and the observations have to be explained using an interacting theory. The fact that one cannot (at least so far) explain these type of observations in a formally satisfying interacting theory of quantum fields is just a point against the QFT framework itself.

The particle objects to which we propose to assign ontological commitment are those that carry interaction charges, and can thus be seen as localised *bundles* of such charges. Following our argument of using interactions as a basis for an ontology of objects in the world, this would constitute the most parsimonious reading of the SM ontology: we commit to the existence to interaction charges, and we posit that particles are gauge-invariant localised bundles of these charges. An electron would be its mass, its spin, and its electroweak charge—indeed the properties with which physicists define what an electron is. As opposed to the fields of the QFT framework, the charges carried by particle states are *actual*, in that we can see how they affect chains of causes and effects in our experimental devices.

This sums up our proposal towards an ontology of the SM. If the framework of QFT is not interpretable in terms of particles, this only shows how inadequate this framework is as a model of fundamental reality. Instead, QFT is best seen as an effective way of describing physical systems within the regime in which special relativistic

⁷ See Chakravarty (2020) for a recent argumentation along these same lines.

⁸ With this we mean the following: the LHC reaches energy levels of $s = 13$ TeV, which, as an order of magnitude, would correspond to exploring distances at the scale of $hc/s \sim 10^{-19}$ meters.

and quantum effects are dominant. The SM, as an interaction theory that models phenomena in such a regime, can be assigned an ontology in terms of particles that undergo collisions, decays, creation and annihilation processes, and the myriad of effects that have been repeatedly tested in particle accelerators.

7 Fundamentality

At this point another relevant critique can be raised: these particle objects could conceivably only be the effect of the collective behaviour of so far unknown underlying fundamental entities. This is very commonly the case in condensed matter systems, where the collective excitations of the underlying atomic lattice allow for a description in terms of ‘emergent’ localised degrees of freedom. Examples include phonons, ‘hole’ quasi-particles in semiconductors (and also the phenomenon of electrons acquiring an effective mass within these materials), Cooper pairs, vortices, etc. As this could also be the case for the measured quantities of the SM, the proposed ontology of fundamental particles could be put in doubt, as these particles would only be a collective effect of as-yet-undetected fundamental components—just as we argued is the case for fields.

Indeed, there are several ideas for a fundamental ontology that go in this direction. Perhaps most famously, although also very speculatively, string theory would purportedly explain all the content of the SM as vibration modes of fundamental one-dimensional objects. In this case, the particles of the SM would be the excitations of these underlying fundamental objects.

We have several answers to this problem, which ultimately lead us to refine our proposal for an ontology of the SM. The first and most obvious is that we do not need to deal with absolute fundamentality when evaluating the ontology of a theory: the SM ontology could perfectly well be a particle ontology, even if this is not the most fundamental theory of matter. In the same way, one could say that the ontology for chemistry is given by atoms, or that the ontology for social science is in terms of persons. However, as scientific realists we would like to say something deeper than that when analysing our most fundamental theories. The interesting question is then: what is the best candidate to a fundamental ontology that we can discern given our current knowledge about fundamental physics, which is indeed mainly encoded in the SM?

The possibility that the SM particles could be non-fundamental does not take away from the fact that they enter into interactions that can be described, predicted, and manipulated—the factors we use as a guide for our ontology in the case of any interaction theory. The particles of the SM are undeniably a part of the furniture of the world, even though we might not be able to tell what are the fundamental building blocks behind them (building blocks that maybe are unobservable even in principle, such as in Esfeld and Deckert (2017)). In this we agree with Falkenburg (2007), when she argues for giving ontological weight to collective excitations in condensed matter—and by extension to the particles of the SM. Notice that an ontology in terms of quantum fields, however, suffers an additional blow from these arguments. The fact that the framework theory for the purportedly non-fundamental entities of the SM is

constructed in terms of such abstract objects as quantum fields should bear very little weight when choosing a fundamental ontology of the world.

In brief, a fundamental ontology of entities from the point of view of the SM is best understood in terms of *localised, gauge-independent* quanta, which are naturally associated with the notion of particle that has proven empirically useful. We can call this an *effective ontology* of particles.⁹ This notion of effective particle can best be understood by referring to the Källén-Lehmann spectral representation (Weinberg, 1996) for the 2-point correlation function of an interacting QFT, or equivalently the so-called ‘propagator’ of the theory. This representation allows us to see the correlation between space time points in a full interacting QFT as an (infinite) sum of particle propagators for a free (interactionless) field theory. This sum is weighted in terms of what is known as the spectral density function, which can be rigorously defined, even in the context of formal algebraic QFT. For a free theory, the spectral density can be seen as a sum over Dirac delta functions of the free particle masses. In an interacting theory, the spectral density will in general be a real positive function of the 4-momentum (more precisely, of its modulus), with the peaks in this function interpreted as one or many particle states (including possible bound states).

The finite size of the peaks of the spectrum density are due to the unstable nature of most particles in interacting theories. The width of the peak is related in a well-defined way to the unstable particle lifetime. Accordingly, the following approach appears valid: to interpret the peaks in the spectral density (which is an experimentally measurable function) as describing the energies and lifetimes of entities which behave as effective particles. These effective particles are very real as carriers of the fundamental charges of the SM, and they could be—as far as we know—the ultimate constituents of reality. It is true that from a Parmenidean point of view there is something bothersome in an ontology of entities that can appear or disappear to and from nothingness (i.e., unstable particles). However, the SM and the associated particle collision experiments show us that there are charges (such as strangeness, or charm) that only manifest in the world for very brief periods of time. Being associated with interaction theories, one should take these charges seriously, and accept that fundamental elements of our ontology can be ephemeral.

We run into some complexities when considering the Higgs mechanism, famously responsible for the masses of most fundamental particles (Friederich, 2014; Benitez et al., 2022). Mass in the SM is not an intrinsic property that works as an interaction charge. Instead, mass is dynamically generated by the interaction of fundamental particles with what is known as the vacuum expectation value of the Higgs field, a scalar field. This field interacts via the electroweak force with quarks and leptons, as it is a (electroweak) charged field. Thus, it would appear that an ontology in terms of fields is unavoidable for this vacuum expectation value, as mass is not generated by interaction with the Higgs *particle*, only by interactions with the Higgs vacuum expectation value field. Notably, even though the quantum Higgs fields is trivially gauge-dependent, its vacuum expectation value is akin to a classical field and is also gauge invariant, so several of our arguments against a field ontology do not apply in this case.

⁹ We thank an anonymous referee for proposing this terminology.

To solve this issue, one option is to consider the Higgs vacuum expectation value—again, a gauge invariant, classical scalar field, not a quantum superposition of fields—as a real substance, which interacts with matter particles. Thus, in the case of the Higgs boson, its quantum Lagrangian would not only model the behaviour of Higgs particles, but also that of the interactions between matter and the vacuum expectation value of the scalar field. A classical field is a much less abstract object than a quantum field, and the ontology of the SM could perfectly well be understood in terms of particles and classical fields. Every particle that has a mass is ultimately fundamentally characterised by a different *charge*: the coupling strength to the Higgs field. In this way, mass could once again be considered a charge-like property.

To summarize, the SM is best served by an ontology of *effective particles*, which are localized quanta of matter, invariant in their description under gauge transformations, and carrying the charges of the four fundamental forces (the three actual charges of the SM, plus the gravitational ‘charge’ of mass-energy). Our ontological commitment to these particles comes mainly from their charge-carrying capacity, as we commit to the existence of the charge properties that ground the interactions we observe in nature. These effective particles can be described by our theories in several ways depending of the context. In the context of Lagrangian QFT, they can be associated with the ‘in’ and ‘out’ free particle states of perturbative calculations, or with the peaks in the the Källén-Lehmann spectral representation of the SM propagator. Within the context of algebraic QFT, effective particles cannot (as of yet) be described, to the detriment of this framework.

We consider this effective particle ontology to be a viable ontology for the standard model of particle physics, one that accounts for our best science, for its provisional character, and for the framework/interaction aspects of it. This proposal avoids many of the issues of previous attempts, taking into account the many reasons why quantum fields are not to be considered part of the actual furniture of the world, without taking away from the impressive success of the SM. In doing so, it allows for an ontology that is not only compatible with scientific observations, but that is parsimonious in terms of metaphysical commitments: the property of interaction charge plays the central role, and particle states can be seen as bundles of such charges. We thus commit to these properties as parts of what is real, a commitment that is justified by the very means we have of knowing the world, namely its interactions. By prioritising in this way an ontology of interactions over that of the theoretical frameworks used in their description, this view is as parsimonious as possible, while still following what our best science tells us about the world.

This proposal goes beyond a naïve realism about SM phenomenology, as it takes into account lessons learned from the study of the renormalisation group and a meta-theory of gauge symmetries. Gauge-independent, localised bundles of interaction charges are a long distance from the classical intuitive image of a matter point, while still avoiding exotic metaphysics such as relationships without relata, Platonic ideals, or highly abstract mathematical objects crowding every point of space-time. In our metaphysically parsimonious view, interactions are what grounds causes and effects, charges are what grounds interactions, and the furniture of the world is constituted by physically instantiated such charges.

To conclude this section, in the future the SM might well not be our most fundamental physical theory. But the ontology of this theory is independent of us believing the theory to be fundamental or not. It is clear following our arguments that the SM as a theory is best understood in terms of particles, and not of quantum fields. If the day comes when theoretical ideas such as super-symmetry or grand unification schemes are neatly worked out and proven to correspond to observations, the symmetry groups that are considered to be fundamental will change radically, as well as the nature of the fields themselves. Yet, even in this scenario, the fundamental particles known to us would be stable enough with respect to such theory change. Equivalently, if it comes to pass that the physicists' community demonstrates that a future complete string theory is in agreement with experiments, and explains the dynamics of all four forces in a unified view, we would be dealing with fundamental objects—(super)strings—which are not described by QFT as we know it. Particles, however, would still be present in such a view, having an existence as vibration modes of the purported fundamental strings. Similar considerations can be made if particles are seen as emerging from collective excitations of underlying, so far undetected, degrees of freedom. In brief, the point is none other than this: what will surely be preserved in a future theory are the particle descriptions, whereas quantum fields will most surely come to be seen simply as an effective ways of describing particle interactions at 'low' energies.

8 Conclusions

In this work we propose an ontology for the standard model of particle physics based on particle-like objects. To justify this we make use of a version of selective scientific realism standing on the notion that there are two kinds of scientific theories. On the one hand there are theories that serve as frameworks, constraining what states of affairs or structures are possible in the world, and on the other hand there are interaction theories, dealing with the observed (and inferred) interactions between entities in the empirically accessible world. This distinction, we contend, must have ontological implications, and provides a useful guide for a selective version of scientific realism.

These ideas lead to a notion of selective realism giving more weight to the ontological status of objects belonging to interaction theories, and even more so in cases where the framework at hand can be considered incomplete or in some other way unsatisfactory. Part of the issue is that, as a scientific realist, one inevitably has to do with non-final, non well-defined, or otherwise incomplete theories. In these contexts, the selective realist should follow the most direct sources of information stemming from science, which is information about interactions, and not the highly indirectly derived information that builds up framework theories.

Here we argue that these notions can be illuminating in the case of the interaction theory of the standard model of particle physics, set within the framework of quantum field theory. As is well-known, even though this is one of our most successful scientific theories, the available framework, based on special relativity and quantum mechanics, cannot formally deal with the notion of point particle when interactions are present, even if particle-like objects empirically seem to play a role in the phenomena we observe.

We emphasise that (i) framework and interaction theories should not always be expected to be fully compatible, as tensions between them have historical precedents; (ii) in general more ontological weight should be given to the objects required by interactions theories—the charge carriers. In our proposal, these objects—which we have dubbed ‘effective particles’—would be identified either with the asymptotic ‘free-particle’ states used within most computations in the SM, or with the quasi-particles in the spectrum of the theory, which can be observed experimentally; (iii) the framework of QFT has been shown, by the formalism of the renormalisation group, to be generally only valid within a certain (eventually large but finite) range in the energy scale, so that QFT is by construction not necessarily amenable to an analysis of ‘the ultimate building blocks of matter’; (iv) any future theory of fundamental physics may dispense with the notion of quantum fields, but will undoubtedly have to deal with explaining the observed properties of particle collisions and decays.

Being defined within the framework of QFT, the SM as is understood today is by necessity a QFT. This, however, only means that up to now the best way we have found to express the interactions between fundamental objects has been within this framework. But we have several reasons to doubt the framework, including the very issue of the impossibility of describing localised states in interacting QFTs. Indeed, the fact that an interacting QFT cannot describe electrons should be seen as an issue for the theory itself, instead of an issue for considering electrons or quarks to be real objects.

Ultimately, what we claim is that metaphysics should not give much ontological weight to framework theories, to which we only have very indirect access. Instead of prioritising the constraints and abstract objects of special relativity and quantum mechanics, we should base an ontology of objects on the interactions we observe in the world. The objects to which we should commit are those that carry in themselves the charges associated with these interactions, interactions which allow us to detect and measure the world, and by means of which we can abstract framing theories about general constraints and structures.

More generally, these ideas themselves can serve as a framework of sorts for future discussions. In particular, we don’t pretend to have exhausted the ontological implications of the framework/interaction theory dichotomy. As this work shows though, the specifics of each theory need to be taken into account when assigning ontological commitments, on top of the framework/interaction distinction.

It would be interesting to extend these notions beyond the SM. In this work we leave largely unexplored a big region of the present-day fundamental physics landscape, that of quantum gravity research, and of the proposed completions of the SM, such as supersymmetry. We think the ideas presented up to now can be applied also to these cases, and the knowledgeable reader can perhaps grasp what would be some of the consequences of taking the point of view here presented to cases such as e.g., string theory. Given their complexity, and the fact that these theories are in many ways very incomplete or provisional, and still not an official part of our scientific knowledge, we leave these questions for future work.

Acknowledgements Thanks to Michael Esfeld and Frida Trotter for a critical reading of the manuscript, and to the anonymous referees that provided ideas to improve it. This work was partially supported by the Swiss SNF, project number 105212-175971.

Funding Open access funding provided by University of Bern

Declarations

Conflicts of interest The author declares no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Aharonov, Y., & Bohm, D. (1959). Significance of electromagnetic potentials in the quantum theory. *Physical Review*, 115, 485–491.

Araki, H., et al. (1999). *Mathematical theory of quantum fields* (Vol. 101). Oxford University Press.

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

Barrett, J. A. (2002). On the nature of measurement records in relativistic quantum field theory. In M. Kuhlmann, H. Lyre, A. Wayne (Eds.), *Ontological aspects of quantum field theory*, (pp. 165–179). World Scientific.

Belot, G. (2003). Symmetry and gauge freedom. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34(2), 189–225.

Benítez, F. (2019). Selective realism and the framework/interaction distinction: A taxonomy of fundamental physical theories. *Foundations of Physics*, 49(7), 700–716.

Benítez, F., Romero-Maltrana, D., & Razeto-Barry, P. (2022). (re) interpreting $E = mc^2$. *Foundations of Physics*, 52(1), 1–19.

Berges, J., Tetradis, N., & Wetterich, C. (2002). Non-perturbative renormalization flow in quantum field theory and statistical physics. *Physics Reports*, 363(4), 223–386. Renormalization group theory in the new millennium. IV.

Berghofer, P. (2018). Ontic structural realism and quantum field theory: Are there intrinsic properties at the most fundamental level of reality? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 62, 176–188.

Bigaj, T. (2018). Are field quanta real objects? some remarks on the ontology of quantum field theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 62, 145–157.

Brading, K., & Castellani, E. (2003). *Symmetries in physics: philosophical reflections*. Cambridge University Press.

Butterfield, J., & Bouatta, N. (2015). Renormalization for philosophers. *Metaphysics in Contemporary Physics*, 104, 437–485.

Cardy, J. (1996). *Scaling and renormalization in statistical physics*, (Vol. 5). Cambridge University Press.

Carroll, B. W., & Ostlie, D. A. (2017). *An introduction to modern astrophysics*. Cambridge University Press.

Chakravarty, A. (2020). Realist representations of particles: The standard model, top-down and bottom-up. Clifton, R., Bub, J., & Halvorson, H. (2003). Characterizing quantum theory in terms of information-theoretic constraints. *Foundations of Physics*, 33(11), 1561–1591.

Costello, K. (2011). *Renormalization and effective field theory* (Vol. 170). American Mathematical Soc.

Dehmelt, H. (1988). A single atomic particle forever floating at rest in free space: New value for electron radius. *Physica Scripta*, 1988(T22), 102.

Delamotte, B. (2002). A hint of renormalization. <https://arXiv.org/0212049>

Delamotte, B. (2012). An introduction to the nonperturbative renormalization group. In A. Schwenk & J. Polonyi (Eds.), *Renormalization group and effective field theory approaches to many-body systems* (pp. 49–132). 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., & Fraser, D. (2006). Haag's theorem and its implications for the foundations of quantum field theory. *Erkenntnis*, 64(3), 305.

Egg, M. (2014). *Scientific realism in particle physics: A causal approach* (Vol. 29). Walter de Gruyter GmbH & Co KG.

Egg, M., Lam, V., & Oldofredi, A. (2017). Particles, cutoffs and inequivalent representations. Fraser and Wallace on quantum field theory. *Foundations of Physics*, 47, 453–466.

Einstein, A. (1919). Time, space, and gravitation. *Times (London)*, pp. 13–14.

Esfeld, M., & Deckert, D.-A. (2017). *A minimalist ontology of the natural world*. Routledge.

Falkenburg, B. (2007). *Particle metaphysics: A critical account of subatomic reality*. Springer.

Flores, F. (1999). Einstein's theory of theories and types of theoretical explanation. *International Studies in the Philosophy of Science*, 13(2), 123–134.

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. (2009). Quantum field theory: Underdetermination, inconsistency, and idealization. *Philosophy of Science*, 76(4), 536–567.

Fraser, D. (2011). How to take particle physics seriously: A further defence of axiomatic quantum field theory. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, 42(2), 126–135.

French, S. (2014). *The structure of the world: Metaphysics and representation*. OUP Oxford.

Friederich, S. (2014). A philosophical look at the Higgs mechanism. *Journal for General Philosophy of Science*, 45(2), 335–350.

Galison, P., et al. (1997). *Image and logic: A material culture of microphysics*. University of Chicago Press.

Georgi, H. (1993). Effective field theory. *Annual review of nuclear and particle science*, 43(1), 209–252.

Glick, D. (2016). The ontology of quantum field theory: Structural realism vindicated? *Studies in History and Philosophy of Science Part A*, 59, 78–86.

Haag, R. (2012). *Local quantum physics: Fields, particles, algebras*. Springer.

Halvorson, H., & Clifton, R. (2002). No place for particles in relativistic quantum theories? In M. Kuhlmann, H. Lyre, & A. Wayne (Eds.), *Ontological aspects of quantum field theory*, (pp. 181–213). World Scientific.

Halvorson, H., & Müger, M. (2006). Algebraic quantum field theory. <https://arXiv.org/0602036>

Healey, R. (2007). *Gauging what's real: The conceptual foundations of contemporary gauge theories*. Oxford University Press.

Jakubczyk, P., Dupuis, N., & Delamotte, B. (2014). Reexamination of the nonperturbative renormalization-group approach to the Kosterlitz-Thouless transition. *Physical Review E*, 90(6), 062105.

Kragh, H. (2002). *Quantum generations: A history of physics in the twentieth century*. Princeton University Press.

Kuhlmann, M. (2013). *The Ultimate Constituents of the material world: In search of an ontology for fundamental physics*, (Vol. 37). Walter de Gruyter.

Kuhlmann, M., & Stöckler, M. (2018). Quantum field theory. In C. Fribe, & M. Kuhlmann (Eds.), *The Philosophy of Quantum Physics*, (pp. 221–262). Springer.

Ladyman, J. (1998). What is structural realism? *Studies in History and Philosophy of Science Part A*, 29(3), 409–424.

Ladyman, J., & Ross, D. (2007). *Every thing must go: Metaphysics naturalized*. Oxford University Press.

Lange, M. (2001). The most famous equation. *The Journal of philosophy*, 98(5), 219–238.

Lange, M. (2007). Laws and meta-laws of nature. *The Harvard Review of Philosophy*, 15(1), 21–36.

Lange, M. (2007). Laws and meta-laws of nature: Conservation laws and symmetries. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(3), 457–481.

Lazarovici, D. (2018). Against fields. *European Journal for Philosophy of Science*, 8(2), 145–170.

Le Bellac, M. (1991). *Quantum and statistical field theory*. Clarendon Press.

Lyre, H. (2004). Holism and structuralism in $U(1)$ gauge theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 35(4), 643–670.

Malament, D. B. (1996). In defense of dogma: Why there cannot be a relativistic quantum mechanics of (localizable) particles. In R. Clifton (Ed.), *Perspectives on quantum reality*, (pp. 1–10). Springer.

Maltrana, D., Herrera, M., & Benítez, F. (2022). Einstein's theory of theories and mechanicism. In J. R. Brown, & M. Carrier (Eds.), *International Studies in the Philosophy of Science*, (pp. 1–18).

Mooser, A., Ulmer, S., Blaum, K., Franke, K., Kracke, H., Leiteritz, C., Quint, W., Rodegheri, C. C., Smorra, C., & Walz, J. (2014). Direct high-precision measurement of the magnetic moment of the proton. *Nature*, 509(7502), 596.

Moretti, V., et al. (2018). *Spectral theory and quantum mechanics: Mathematical foundations of quantum theories. Symmetries and introduction to the algebraic formulation*. Springer.

Oldofredi, A. (2018). No-go theorems and the foundations of quantum physics. *Journal for General Philosophy of Science*, 49(3), 355–370.

Paul, L. A. (2017). *A one category ontology*. Yale University.

Redhead, M. (2002). The interpretation of gauge symmetry. In M. Kuhlmann, & H. Lyre (Eds.), *Ontological aspects of quantum field theory*, (pp. 281–301). World Scientific.

Romero-Maltrana, D., Benítez, F., & Soto, C. (2018). A proposal for a coherent ontology of fundamental entities. *Foundations of Science*, 23, 705.

Ruetsche, L. (2002). Interpreting quantum field theory. *Philosophy of Science*, 69(2), 348–378.

Ruetsche, L. (2012). Philosophical aspects of quantum field theory: I. *Philosophy Compass*, 7(8), 559–570.

Summers, S. J. (2012). *A perspective on constructive quantum field theory*. <https://arxiv.org/1203.3991>

Vaidman, L. (2012). Role of potentials in the Aharonov-Bohm effect. *Physical Review A*, 86(4), 040101.

Weatherall, J. O. (2015). *Categories and the foundations of classical field theories*. <https://arxiv.org/1505.07084>

Weatherall, J. O. (2016). Understanding gauge. *Philosophy of Science*, 83(5), 1039–1049.

Weinberg, S. (1995). *The Quantum Theory of Fields*, (Vol. 1). Cambridge University Press.

Weinberg, S. (1996). *The Quantum Theory of Fields*, (Vol. 2). Cambridge University Press.

Williams, P. (2019). Scientific realism made effective. *The British Journal for the Philosophy of Science*, 70(1), 209–237.

Wilson, K. G. (1975). The renormalization group: Critical phenomena and the Kondo problem. *Reviews of modern physics*, 47(4), 773.

Wilson, K. G., & Fisher, M. E. (1972). Critical exponents in 3.99 dimensions. *Physical Review Letters*, 28(4), 240.

Wilson, K. G., & Kogut, J. (1974). The renormalization group and the ϵ expansion. *Physics Reports*, 12(2), 75–199.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.