In defence of naiveté: The conceptual status of Lagrangian quantum field theory

David Wallace

December 23, 2001

Abstract

I analyse the conceptual and mathematical foundations of Lagrangian quantum field theory (that is, the 'naive' quantum field theory used in mainstream physics, as opposed to algebraic quantum field theory). The objective is to see whether Lagrangian quantum field theory has a sufficiently firm conceptual and mathematical basis to be a legitimate object of foundational study, or whether it is too ill-defined. The analysis covers renormalisation and infinities, inequivalent representations, and the concept of localised states; the conclusion is that Lagrangian QFT (at least as described here) is a perfectly respectable physical theory, albeit somewhat different in certain respects from most of those studied in foundational work.

1 Introduction

From its beginning, quantum field theory (QFT) has been plagued with mathematical difficulties. Any attempt to apply the theory to interacting systems led to the appearance of infinities, and although methods were found[1] to remove those infinities from the calculated cross-sections by cutting off certain integrals at finite (though very high) energies, these methods had a sufficiently *ad hoc* feel that the theory was not felt to be on remotely satisfactory conceptual ground.

From the 1950s onwards, there has been a major program to resolve this problem by reformulating QFT on an axiomatic basis: that is, starting from what seem to be physically necessary — and mathematically precise — principles which any QFT would have to satisfy, and then finding QFTs which actually satisfy them. This program is now generally referred to as *algebraic quantum field theory*, or AQFT; see Haag (1996), and references therein, for extensive discussion of it.

[1] By Dyson, Feynmann, Schwinger and Tomonaga, amongst others; see Cao (1997, pp. 185–209) for a historical discussion.

The major problem with AQFT is that very few concrete theories have been found which satisfy the AQFT axioms. To be precise, the only known theories which do satisfy the axioms are interaction-free: no examples are known of AQFT-compatible interacting field theories, and in particular the standard model cannot at present be made AQFT-compatible.

Nonetheless, the great majority of foundational discussions of QFT are carried out in an AQFT framework, for obvious reasons: the precise axiomatic nature of AQFT makes it ideal for foundational study, since we can specify with mathematical precision exactly what the entity is that we are studying.

Meanwhile, the mainstream physics community has continued to work with what we might call (for want of a better name) "Lagrangian" QFT, the original QFT developed in the 1930s and whose infinities were tamed (at least pragmatically speaking) in the 1950s. In Lagrangian QFT, an entirely different program exists for understanding those infinities, in terms of the renormalisation group; this program, due primarily to Wilson, does not eliminate the finite cutoffs introduced to deal with the infinities, but rather attempts to 'legitimize' them. It is in the framework of Lagrangian QFT that the Standard Model is formulated.

This paper is an investigation of whether Lagrangian QFT is sufficiently well-defined conceptually and mathematically that it too can be usefully subjected to foundational analysis. The reasons for making such an investigation are threefold. Firstly, the problem with restricting our foundational studies to AQFT is that — pending the discovery of a realistic interacting AQFT — we have only limited reason to trust that our results apply to the actual world, which appears to be described rather well by the Standard Model.

One response to this criticism is to say that Lagrangian QFT is so mathematically ill-defined that we cannot regard it as a proper theory at all, so we cannot trust its results either! Our second reason for the investigation is then to see whether this criticism is justified.

The third motivation is somewhat more philosophical. Wilson's explanation of the renormalisation procedure relies upon the failure of the QFT to which it is applied at very short distances. It is then intriguing to ask how to put on a firm conceptual footing a theory which relies for its mathematical consistency on its own eventual failure.

The structure of the paper is as follows. Section 2 reviews the formal definition of a Lagrangian QFT. This definition is very formal and mathematically totally ill-defined; sections 3 and 4 address this problem, dealing respectively with renormalisation and with the existence of inequivalent representations of the observable algebra. Section 5 considers the problem of how to define local objects in QFT. This is also a problem in AQFT and much of this section's argument applies in that domain; however, the framework of sections 2–4 offers some interesting insights. Section 6 is the conclusion.

Three comments should be made before the main part of the paper:

1. This paper is not intended as an attack on the AQFT program. The program to construct an AQFT-compatible interacting field theory is well-motivated and, if successful, would be of enormous import; the foundational results which have emerged from AQFT have been of considerable importance in understanding quantum field theory and in general they apply also to Lagrangian QFTs. This paper should be read as complementary to, rather than in competition with, these results.

2. There is a strong tradition in foundational studies of QFT to treat the theory from an operational viewpoint, formalised by regarding the local algebras of AQFT as describing operations which can be carried out by an observer localised in a given region. This paper does not follow that viewpoint: it treats QFT instead as a closed theory (a 'universal theory' in Deutsch's (1985) sense), and regards measurement devices, observers etc. as just subsystems of the QFT state. This difference of approach has only limited significance for most of the paper but is important in section 5.3.

3. Although much of the discussion below applies equally to bosonic and fermionic field theories, the specific examples discussed are exclusively bosonic.

2 What is a Lagrangian QFT?

For better or for worse, most canonical quantum field theories are found by starting with a classical field theory and then 'quantizing' it. To be sure, there is something intellectually unsatisfactory about this: given that quantum theory is the more fundamental theory, we would prefer to work in the other direction, that is, to recover *classical* field theories from quantum starting points (see Deutsch (1984) for a development of this criticism). Nonetheless, the classical starting point has proven a powerful method for finding QFTs, and we adopt it here.

2.1 Classical field theories

A classical relativistic field theory can be considered as consisting of a relativistic spacetime $\mathcal{M}$ (such as Minkowski spacetime), a set of fields (that is, maps from $\mathcal{M}$ to some other space, such as the real or complex numbers[2]), and a Lagrangian density: a real-valued function $\mathcal{L}$ on $\mathcal{M}$ whose value at a given point x depends only on the fields and their first partial derivatives, evaluated at x. For a given region $\mathcal{D}$ of $\mathcal{M}$, we define the action $S_{\mathcal{D}}$ as the integral of $\mathcal{L}$ over $\mathcal{D}$; $S_{\mathcal{D}}$ is thus a functional of the fields, and we define

[2] In more advanced treatments, we might take the fields to be sections of some fibre bundle over $\mathcal{M}$; this technicality does not significantly affect the arguments of the paper, and will be ignored.

the dynamically allowed field configurations within $\mathcal{D}$ to be those for which $S_{\mathcal{D}}$ is extremal under variations of the fields which vanish at the boundary of $\mathcal{D}$.

To go further, it is necessary to take $\mathcal{M}$ to be globally hyperbolic, and then to fix a foliation $\mathcal{M} = \Sigma \times \mathrm{R}$. Let Σ_1, Σ_2 be any two spatial slices in the foliation, with associated time coordinates t_1, t_2 respectively; let the spacetime region between the two slices be $\mathcal{D}_{12}$. If the fields are required to fall off rapidly enough at spatial infinity, then we can define the action $S_{\mathcal{D}_{12}}$:

$$S_{\mathcal{D}_{12}} = \int_{\mathcal{D}_{12}} \mathrm{d}^4\mu \, \mathcal{L} \; = \int_{t_1}^{t_2} \mathrm{d}t \, \int_{\Sigma} \mathrm{d}^3\mu(t) \mathcal{L} \tag{1}$$

where $\mathrm{d}^4\mu$ is the four-dimensional volume element on $\mathcal{M}$ and $\mathrm{d}^3\mu(t)$ is the induced three-dimensional volume element on Σ. As the notation indicates, the three-volume element is in general time-dependent; however, when $\mathcal{M}$ has a time-translation symmetry then we can always choose a foliation such that this time- dependence vanishes.

We can now define the Lagrangian $L(t)$ of the theory by

$$L(t) = \int_{\Sigma} \mathrm{d}^3\mu(t) \mathcal{L}. \tag{2}$$

L is now a functional of the fields on Σ and their time-derivatives, and the action $S_{\mathcal{D}_{12}}$ is given simply by

$$S_{\mathcal{D}_{12}} = \int_{t_1}^{t_2} \mathrm{d}t \, L(t), \tag{3}$$

so the field theory is now in Lagrangian form (albeit infinite-dimensional).

The final step before quantisation is to transform from Lagrangian to Hamiltonian form, which we do by a straightforward infinite-dimensional generalisation of the Legendre transform in classical mechanics: to each field ϕ_i we associate a conjugate momentum field π_i by

$$\pi_i(\mathbf{x}) = \frac{\delta L}{\delta \dot{\phi}_i(\mathbf{x})}, \tag{4}$$

and we define the Hamiltonian H by

$$H[\phi, \pi] = \sum_i \int_{\Sigma} \pi_i(\mathbf{x}) \dot{\phi}_i(\mathbf{x}) - L[\phi, \dot{\phi}], \tag{5}$$

where $\dot{\phi}_i$ is defined implicitly in terms of the π_i, via (4). Points in the phase space $\mathcal{P}$ of the field are then given by specifying sets of functions (ϕ_i, π_i) (for all i); it is easy to check that the dynamics are given by Hamilton's equations, as in the finite-dimensional case:

$$\dot{\phi}_i(\mathbf{x}) = \frac{\delta H}{\delta \pi_i(\mathbf{x})}; \tag{6}$$

$$\dot{\pi}_i(\mathbf{x}) = -\frac{\delta H}{\delta \phi_i(\mathbf{x})}. \tag{7}$$

The Poisson bracket on $\mathcal{P}$ is given by

$$\{A[\phi,\pi], B[\phi,\pi]\} = \sum_i \int_\Sigma \mathrm{d}^3\mathbf{x} \left(\frac{\delta A}{\delta\phi_i(\mathbf{x})}\frac{\delta B}{\delta\pi_i(\mathbf{x})} - \frac{\delta A}{\delta\pi_i(\mathbf{x})}\frac{\delta B}{\delta\phi_i(\mathbf{x})} \right), \tag{8}$$

and ϕ_i and π_i obey natural infinite-dimensional analogues of the canonical relations: $\{\phi_i(\mathbf{x}), \phi_j(\mathbf{y})\} = \{\pi_i(\mathbf{x}), \pi_j(\mathbf{y})\} = 0$ and $\{\phi_i(\mathbf{x}), \pi_j(\mathbf{y})\} = \delta_{i,j}\delta(\mathbf{x}-\mathbf{y})$.

Before going on to quantise this theory, we make two observations:

1. Despite the somewhat casual approach of this section, classical field theories have perfectly well-defined Lagrangian and Hamiltonian forms, and the above analysis can be carried out in a fully rigorous way; see chapter 7 of Woodhouse (1991), chapter 3 of Marsden and Ratiu (1994) and references therein for such an analysis. In this more rigorous approach it is necessary to replace the point functionals $\phi_i(\mathbf{x})$ and $\pi_i(\mathbf{x})$ with smoother functionals on phase space, and to eschew use of the delta-function Poisson bracket between such point functionals; such objects have a status not dissimilar to that of position eigenstates in non-relativistic quantum mechanics, in that they cannot easily be rigorously defined but (if used with a little caution) are of great conceptual and calculational convenience.

2. The almost immediate introduction of a foliation of $\mathcal{M}$ may cause some readers to worry about violation of relativistic covariance. However, recall that each point in $\mathcal{P}$ defines a unique trajectory through $\mathcal{P}$, so that we may set up an isomorphism between phase-space and the space of all solutions to the field equations. The latter space *is* covariant (in the sense that it can be defined without use of any preferred foliation) and if desired all talk of 'phase space' may be reinterpreted as talk of solution space.

Both of these points are developed further in Wald (1994) (from an AQFT perspective, and in the somewhat restricted context of free fields).

For convenience, in the remainder of this paper we assume that we are dealing with a single real field, and so drop the subscripts on ϕ and π; this has no significant consequences.

2.2 Quantization: observables

In an ideal world, quantization would work as follows: it would yield a Hilbert space $\mathcal{H}$, together with a map Q from the functions on $\mathcal{P}$ (which can be thought of as the classical observables) into the space of self-adjoint operators on $\mathcal{H}$. Q would have the property that, if $\{,\}$ is the classical Poisson bracket, then $Q(\{A,B\} = i(Q(A)Q(B) - Q(B)Q(A))$.

This is not an ideal world: quantization as defined above is provably impossible (see, e. g. , Abraham and Marsden (1978) for a discussion). What is possible is to find such a Q for a much more restricted class of observables: in nonrelativistic particle mechanics, for instance, we choose spatial position and the momentum conjugate to it. There is no known algorithm for choosing this restricted class, and no real reason to expect one to exist (c. f. the criticisms of quantization in Deutsch 1984); two different choices will sometimes lead to empirically inequivalent quantum theories, in which case experiment is the only way to determine which is correct.[3] In general, however, Nature is fairly kind to us, and making the most obvious choice of observables tends to work. QFT is no exception: choosing $\phi(\mathbf{x})$ and $\pi(\mathbf{x})$ as preferred observables leads to the empirically correct theory.

The question of which observables count as 'fundamental' is also relevant for the interpretation of the quantum theory which we hope to produce. For any fixed dimension, all Hilbert spaces are isomorphic, so it is through the observables that we are able to give physical meaning to a theory. In practice, this usually involves a connection between certain observables and the spacetime of the theory: in NRQM, for instance, it is the designation of the momentum operators as generators of translations which allows us to identify certain states as eigenstates of position.[4] The obvious strategy in QFT is to use the spatiotemporal dependence of the observables to give the theory physical meaning: when quantised, $\phi(\mathbf{x})$ and $\pi(\mathbf{x})$ will become operators which we will treat as localised at $\mathbf{x}$.

This approach can be taken further. First we shift from the Schrödinger to the Heisenberg picture, so that the field operators are functions of spacetime points x rather than just spatial points $\mathbf{x}$. Then, given any (bounded, open) subset $\mathcal{O}$ of $\mathcal{M}$, we define the algebra $\mathcal{A}(\mathcal{O})$ as the algebra of all operators which can be constructed from $\widehat{\phi}(x)$ and $\widehat{\pi}(x)$ whenever $x \in \mathcal{O}$ — so $\mathcal{O}$ includes such operators as $\widehat{\pi}(x)^2 (x \in \mathcal{O})$, $\int_{\mathcal{O}} f(x)\widehat{\phi}(x)$, etc. This gives us (formally) what is technically known as a *net* of operator algebras: a map from bounded, open subsets of a topological space into operator algebras, such that if $\mathcal{O}_1 \subseteq \mathcal{O}_2$, then $\mathcal{A}(\mathcal{O}_1)$ is a subalgebra of $\mathcal{A}(\mathcal{O}_2)$. It is a central claim of algebraic quantum field theory — and one which appears consistent with the way QFT is used in physics — that a QFT is entirely specified by the structure of this algebra net: in other words, that all physical facts about the theory can be determined from knowing the spacetime

[3]Carlip (1998) points out (p. 97) that 'old quantum theory' can be (somewhat anachronistically) understood as wrongly selecting action-angle variables as the preferred observables.

[4]The argument, adapted from Newton and Wigner (1949), is as follows (in one dimension): if a state $|\phi\rangle$ is localised at some point, then any translation $\exp(i\lambda\widehat{P})$ applied to it will leave it localised at a different point; thus $\exp(i\lambda\widehat{P})\,|\phi\rangle$ must be orthogonal to $|\phi\rangle$; hence the position eigenstates are precisely those satisfying $\langle\phi|\exp(i\lambda\widehat{P})\,|\phi\rangle$ for all λ. If $\{|k\rangle\}$ are the (improper) eigenstates of $\widehat{P}$, and if we define $|x\rangle = \int \mathrm{d}k\, \exp(-ikx)\,|k\rangle$, then $\langle x|y\rangle = \delta(x-y)$, and $\exp(i\lambda\widehat{P})\,|x\rangle = |x-\lambda\rangle$; hence, $\langle x|\exp(i\lambda\widehat{P})\,|x\rangle = \langle x|x-\lambda\rangle$, and so the $|x\rangle$ states are exactly localised.

dependence of the observables, without any need to know (for instance) which observable is $\widehat{\pi}(x)$ and which is $\widehat{\phi}(x)$. In particular, this allows the possibility that field theories generated from different classical Lagrangians are actually the same QFT; two such field theories are generally referred to as *Borchers equivalent*.

As such, to specify an algebraic quantum field theory, all that is needed is a net of algebras, which is required to satisfy certain axioms (basically relating to locality and causality; see Haag (1996) for an extended discussion of these axioms). Any such AQFT implicitly specifies a Borchers equivalence class of 'concrete' QFTs each of which generates the same algebra net; all such concrete QFTs are Borchers equivalent to one another, and they are usually regarded simply as different ways of describing the AQFT, analogous to different coordinate systems on a manifold.

The assumption that the net of algebras captures all physical information about a theory gives a very elegant realisation of the idea that quantum theories gain physical content via a specification of the spatiotemporal properties of their observables. We will refer to it as the Net Assumption.

We have glossed one question, though: what does it mean to say that an operator is localised somewhere? In much of the AQFT literature an operationalist answer is given to this question: operators represent physical operations which can be performed, by the observer, on the QFT state, and in particular $\mathcal{A}(\mathcal{O})$ is the algebra of all operations which can be performed, and all measurements which can be made, by an observer who is external to the field system but localised within $\mathcal{O}$. However — as mentioned in the Introduction — we wish to understand a given QFT as self-contained (in which case observers, measurement devices and the like must be built out of states of the QFT) so this approach is not available to us. Instead, we take the expectation values of operators in a given spatial region as giving us information about the degree of excitation of the field in that region. This view will be further developed in the next section.

2.3 Quantization: states

We now address the practical task of actually finding the operator representations of the field observables. In NRQM we usually do this via wave-functions: we take our Hilbert space to be the space of complex functions $\psi(q)$ on the configuration space of the classical theory, and quantize the classical observables q, p via

$$\widehat{q}\psi = q\psi(q); \;\; \widehat{p}\psi = -i\frac{\mathrm{d}\psi}{\mathrm{d}q}. \tag{9}$$

Formally (and only formally; but see sections 3 and 4) we can easily generalise this to infinite-dimensional systems: the configuration space of the system is the infinite-dimensional space of all field configurations on Σ, so the quantum states will be complex-valued functionals on that space. We will denote the Hilbert space of such functionals by

$\mathcal{H}_\Sigma$, and proceed to quantize the classical observables $\phi(\mathbf{x})$ and $\pi(\mathbf{x})$ as

$$(\widehat{\phi(\mathbf{x})}\Psi)[\chi] = \chi(\mathbf{x})\Psi[\chi]; \tag{10}$$

$$(\widehat{\pi(\mathbf{x})}\Psi)[\chi] = -i\frac{\delta\Psi}{\delta\chi(\mathbf{x})}[\chi], \tag{11}$$

where χ is any field configuration on Σ. It is easy to check that infinite-dimensional generalisations of the canonical commutation relations are satisfied:

$$\left[\widehat{\phi(\mathbf{x})}, \widehat{\phi(\mathbf{y})}\right] = \left[\widehat{\pi(\mathbf{x})}, \widehat{\pi(\mathbf{y})}\right] = 0; \tag{12}$$

$$\left[\widehat{\phi(\mathbf{x})}, \widehat{\phi(\mathbf{y})}\right] = i\delta(\mathbf{x} - \mathbf{y}). \tag{13}$$

(For simplicity we have considered only the case of one real field, but the generalisation to multiple fields is obvious.)

Now the Hilbert space $\mathcal{H}_\Sigma$ has a natural tensor-product structure. Let $\Sigma_1, \ldots, \Sigma_n$ be disjoint subsets of Σ whose union is Σ; then (formally) we can define $\mathcal{H}_j$ as the space of complex functionals on functions on Σ_j. Any function f on Σ can be specified uniquely by giving those n functions which are the restrictions of f to each Σ_j; hence we have

$$\mathcal{H}_\Sigma = \otimes_{i=1}^n \mathcal{H}_{\Sigma_i}. \tag{14}$$

If $\mathbf{x} \in \Sigma_j$ then $\widehat{\phi(\mathbf{x})}$ and $\widehat{\pi(\mathbf{x})}$ act trivially on $\mathcal{H}_{\Sigma_i}$ for $i \neq j$. Hence, the $\mathcal{H}_{\Sigma_i}$ can be thought of as representing those subsystems of the field which are localised within Σ_i. This viewpoint can be reversed: we can take spatially localised subsystems of the field as our starting point, in which case $\widehat{\phi(\mathbf{x})}$ and $\widehat{\pi(\mathbf{x})}$ are localised in Σ_j *because* they act trivially on the other $\mathcal{H}_{\Sigma_i}$.

The reader is likely to have noticed that none of the above would win prizes for mathematical rigor! Indeed, on the face of it the approach of this whole section is appallingly badly defined. It will be the purpose of sections 3 and 4 to explain how we can attain peaceful coexistence with this ill-definedness.

3 Infinities and renormalisation

We now have at least a formal definition of a quantum field theory; but it is, clearly, profoundly unsatisfactory in many respects. In this section we consider the question of how, and to what extent, it is possible to restore some degree of mathematical rigour to the theory.

3.1 The problem of infinities

Where to start in analysing the mathematical shortcomings of section 2.3? We could begin by noting the various ambiguities glossed over when considering functionals on an infinite-dimensional space: what topology should be placed on the space of functionals? How smooth must a field configuration be to be admitted?, etc. However, there is a more crucial problem which applies to any attempt to make wave-functional space a Hilbert space. This is the problem of how to define the inner product. In the finite case, of course, we define the inner product of two wave-functions by

$$(\psi_1, \psi_2) = \int_{\Re^n} \mathrm{d}^n q\, \psi_1^*(q_1, \ldots, q_n)\psi_2(q_1, \ldots, q_n); \tag{15}$$

generalising this to wave-functionals gives a functional integral:

$$(\Psi_1, \Psi_2) = \int_{\mathcal{S}} \mathcal{D}\chi \Psi_1^*[\chi]\Psi_2[\chi], \tag{16}$$

where $\mathcal{S}$ is field configuration space and $\mathcal{D}\chi$ is a 'functional measure' on $\mathcal{S}$.

Unfortunately, defining measures on infinite-dimensional spaces is extremely hard. Indeed, simple attempts to define $\mathcal{D}\chi$ tend to give infinity as the result of the functional integral, which clearly isn't satisfactory (see, e. g. , Binney *et al* (1992, p. 409) for an example of how this occurs). The integral can be made well-defined if we restrict ourselves to a finite-dimensional subspace of $\mathcal{S}$, but as soon as we start to consider functions which can vary on arbitrarily short length-scales, we lose the ability to define it.

Blithely ignoring this little problem, we can press on with the development of our QFT, but we soon run into other infinities: in particular, when we try to calculate the effects of interactions, we find that our calculations include terms which involve integrating over arbitrarily short length-scales, and that some of these terms are infinite when the integral is taken over such length-scales. With care we can avoid the infinities in free-field theories (such theories can be exactly defined in AQFT) but no examples are known of realistic QFTs in which these infinities do not have to be confronted.

One highly principled attitude to the problem might be to say: very well, none of our supposed 'interacting QFTs' count as real theories, so let us reject all of them and go looking for properly defined ones. Needless to say, though, this is not the mainstream approach in particle physics, where algorithms for extracting useful information from QFT despite the infinities have been known for nearly fifty years. To understand why these algorithms work — and why they can be understood not just as 'algorithms', but as a valid resolution of the problem of infinities — we take a brief digression into condensed-matter physics.

3.2 High-energy cutoffs and renormalisation

Functional integrals are not restricted to relativistic quantum theory: they occur throughout physics, and in particular in condensed-matter physics. There too they formally lead to infinities; however, there is no question of a conceptual problem. For the integrals are taken on the assumption that (for instance) the density of matter is continuously varying, and can vary on arbitrarily short length-scales. But this continuum assumption can only be an approximation, for matter is made of atoms, and any variations on a length-scale short relative to the interatomic distance clearly lie beyond the scope of the approximation. So the functional integral must be cut off at some short, but finite, length-scale.

It might be thought that this would lead to enormous computational difficulties, for the functional integral is dominated by the short-distance variations and hence the precise details of the cutoff procedure ought to be important. Remarkably, this is not so: it can be shown that if we restrict our attention to the behaviour of the system on lengthscales which are very long relative to the cutoff, then:

1. All but a finite number of possible interaction terms (the so-called 'renormalisable' terms) have negligible effect on the system's behaviour.

2. The only effect of the high-energy (i. e. short-distance) degrees of freedom on the system's long-distance behaviour is to modify ('renormalise') the coefficients of the renormalisable interaction terms. As such, the only effect of getting the cutoff details wrong is to change the effective values of the coefficients in the Hamiltonian. If, instead of trying to calculate these effective values based on the true values (which in any case is often impossible) we simply measure them, the details of the cutoff procedure are completely irrelevant to the long-distance behaviour of the system.

These results were established in the 1970s, primarily by Kenneth Wilson (see Wilson and Kogut (1974) and references therein); Binney *et al* (1992) give an exceptionally lucid exposition.

At least mathematically, this process can be applied equally well to relativistic QFTs: instead of allowing our functional integrals to range over all field configurations, we restrict them to only those configurations whose short-distance variation is not too quick (the conceptually simplest way, at least in Minkowski spacetime, is to exclude all configurations whose Fourier coefficients vanish above some fixed value of $|\mathbf{k}|$, where $1/|\mathbf{k}|$ is the intended cutoff length). This yields a well-defined theory, and if we require the cutoff length to be far smaller than the lengthscales at which we study the theory, then that cutoff length affects the theory's predictions only through renormalisation of the coefficients in the effective Hamiltonian.[5] Since the latter coefficients are in any case only known empirically, this

[5]See chapter 12 of Peskin and Schroeder (1995) for the technical details of this process.

has no practical consequences.

What physical justification might there be for imposing a cutoff in relativistic QFT? Three possibilities are generally mentioned in the literature (see, e. g., Cao (1997) or Binney *et al* (1992) for details):

- It is easy to construct field theories in which some of the degrees of freedom are 'frozen out', i. e. become irrelevant, below some given energy scale. Below this scale, the field theory can be described by an 'effective field theory' (EFT), which does not explicitly include the frozen-out degrees of freedom; however, those degrees of freedom do have the effect of imposing an effective cutoff of the EFT. This is believed to happen for a number of actually studied low-energy theories, whose cutoff energies are experimentally accessible; see Binney *et al* (1992, pp. 372–373), for a brief discussion. Of course, when a cutoff is generated this way, it relies on the existence of the 'true' field theory, which in turn can only be rigorously defined with a cutoff of its own. This implies either some infinite tower of EFTs, which describe physics at successively higher energy scales, or else some other cutoff mechanism which truncates the tower.

- Possibly at sufficiently high energy scales (say, beyond the 'Grand Unification energy' (Peskin and Schroeder 1995), about 10^{16} GeV, above which the strong and electroweak forces are conjectured to become unified) the entire field-theoretic description of physics may break down and be replaced by some other theory (such as string theory). This new theory would then impose an effective cutoff, and might itself be free of infinities.

- At still higher energy levels (and hence still shorter lengthscales), it is widely believed that the concept of spacetime as a continuum will itself break down, to be replaced by some quantized version. On dimensional grounds this is expected to happen at around the Planck energy, or 10^{19} GeV; this is equivalent to a breakdown on lengthscales of order 10^{-34} metres (the Planck length).

In practice, the second and third alternatives may not be all that distinct: although string theory is usually formulated as a perturbative theory on a flat background spacetime, it is generally accepted that this background spacetime must eventually be eliminated from the theory — in which case, presumably, we would again expect spacetime to be an effective concept emerging only on lengthscales greater than the Planck length.

3.3 The conceptual status of a cutoff QFT

The introduction of a finite cutoff, then, mathematically resolves the infinities problem, and might turn out to be justifiable on a number of physical grounds. However, it is prima facie unfortunate from a foundational perspective, for the following reason: in foundational work, it is usual to start with a well-defined (and, preferably, axiomatisable) theory, and then to investigate the implications of that theory. If QFTs are intrinsically approximate theories, which can be trusted only at certain energy scales and whose behaviour is unknown at other scales, it is hard to take this attitude to them. Should we then regard a QFT as a well-defined theory at all, or (as Buchholz (2000) has advocated) as

> an efficient algorithm for the theoretical treatment of certain specific problems in high-energy physics (Buchholz 2000, pp. 1–2)

Cao (1997, pp. 350–352) identifies three different, commonly adopted, attitudes to the foundations of QFT:

(1) The current situation is genuinely unsatisfactory: we should reject the cutoff theories as not mathematically well-defined, and continue to look for nontrivial theories defined at all lengthscales.

(2) The picture of 'an infinite tower of effective field theories' mentioned in the last section is to be taken seriously. As Cao stresses, this would require

> a drastic change of our conception of fundamental physics itself, a change from aiming at a fundamental theory (as the foundation of physics) to having effective theories valid at various energy scales. (Cao 1997, p. 351)

(3) QFTs as a whole are to be regarded only as approximate descriptions of some as-yet-unknown deeper theory, which gives a mathematically self-contained description of the short-distance physics.

(1) is of course the viewpoint of the entire algebraic QFT program: that program's ultimate long-term goal is to produce a mathematically rigorous description of interacting QFTs (specifically, of the Standard Model or of some successor to it). Obviously, completion of that goal would dramatically change the foundational status of quantum field theory; it is only because the goal has not been achieved that it is interesting to investigate alternatives to (1). Conversely, anyone regarding either (2) or (3) as fully satisfactory from a foundational viewpoint would have reason to doubt whether the eventual success of the AQFT program is possible.

(2) will not be discussed further here: adopting it would require dramatic changes to many aspects of our attitude to the foundations of physics, and a discussion of such changes lies far beyond this paper. See Castellani (2000) and Hartmann (2000) for further discussion.

The rest of this discussion will focus on (3), and its attendant foundational problem: how can we give a clean conceptual description of a theory which can be rigorously defined only as the low-energy limit of another theory which we do not yet have?

To see the problem here, contrast this case with a more straightforward one: the current status of classical mechanics (CM), given that we believe it to be in some sense superseded by non-relativistic quantum mechanics (NRQM). What is our justification for continuing — as in many situations we do continue — to use CM, given that we believe NRQM to be the more valid description? It is not enough to observe that NRQM tends to CM as $\hbar \to 0$, for $\hbar$ does not tend to zero: it is a constant. Furthermore, simply considering successively larger systems (so that the quantity ($\hbar$/typical action) tends to zero) is singular and contains various surprises (this matter is discussed in Zurek 1998b). A better approach is to identify what might be called 'classical domains' of NRQM: that is, domains of NRQM in which NRQM is approximately isomorphic to CM. (A system of large bodies decohered by interactions with an environment would be an example of such a domain.) Within such a domain, we could then tentatively apply CM, whilst remaining aware that CM could fail to give accurate answers to some questions, and that its predictions should only be treated as accurate to within the accuracy of the 'approximate isomorphism' between CM and the restricted domain of QM. (This approach to inter-theoretic reduction is discussed in much more detail, and in the specific context of the relationship between CM and NRQM, in Wallace (2001c); it is closely related to Cao's structural approach to QFT (Cao 1997, chapter 12) and to the 'structural realist' program inaugurated by Worrall (1989) and discussed in, e. g. , (Psillos 1995; Ladyman 1998).)

Approach (3) presumes the existence of some theory — call it X — to which QFT is to be an approximation (in some domains). There are, however, a number of important differences between the X–QFT relationship and the NRQM–CM one:

1. Most obviously, we know NRQM, whereas we have only the sketchiest idea of what X will turn out to be. This makes it difficult to delineate the domains in which QFT is approximately isomorphic to X, since we have to describe those domains in QFT-language rather than X-language. In particular, we would like to say that these domains are those in which the states do not vary on too-short lengthscales, but this implicitly suggests that it even makes sense to talk about states which do vary on such lengthscales – which may not be possible in X: indeed, X may not contain any elements which are even approximately isomorphic to too-rapidly-varying states of QFT. In solid-state physics, by analogy, it isn't even meaningful to talk about density fluctuations on lengthscales shorter than the interatomic separation.

2. When using CM, in situations where we are unsure about its validity we can always cross-check with NRQM to check that its predictions agree with CM. Since we don't have X, we can't do this in QFT: the only tests available for its validity in describing a given phenomenon are empirical.

3. CM is a mathematically well-defined and self-consistent theory, whereas QFT as we

have described it so far is not.

It is useful to separate the first and second points from the third. For *if* we had a mathematically well-defined QFT, we would have a perfectly good physical theory which, however, we would have good reason to believe empirically inadequate for describing certain phenomena, even though we have no better alternative. This has been a common situation throughout physics: classical mechanics in the early 20th century is an example of such a theory; general relativity today is another. In fact, until and unless we come up with a theory which we are supremely confident is 'final', this should be our attitude to any well-tested physical theory: its predictive utility and explanatory power require us to take seriously those entities which the theory claims to exist (such as classical particles, or general-relativistic spacetimes), but to remember that

- these entities may be patterns or structures in some more fundamental ontology, rather than fundamental entities in their own right;
- our theory may describe the world not by virtue of being 'fundamental', but by virtue of being instantiated in the structure of a deeper theory;
- this instantiation may be only approximate, leading to domains in which the theory fails.

(See Wallace (2001b) for a discussion of the ontology of patterns and structures in physics.)

So, if we had a mathematically well-defined QFT then we could happily interpret it in the same way we now interpret CM. Can we obtain such a QFT?

The easiest way to do so is to choose a concrete implementation of the short-distance cutoff. A simple recipe for such an implementation is as follows:

1. Choose a cutoff lengthscale l.
2. Choose a spacetime foliation.
3. On each leaf of the foliation, replace the continuum of field operators $\widehat{\phi}(\mathbf{x})$, $\widehat{\pi}(\mathbf{x})$ with a discrete grid of such operators (to be called grid observables) with the grid spacing being $\sim l$.
4. Discretise the Hamiltonian, replacing its integral over functions of field operators with a discrete sum over operators on the grid.
5. Replace the (formal) Hilbert space $\mathcal{H}_\Sigma$ of the QFT (defined in section 2.3 as the space of complex functionals over all field configurations on Σ) with the space of functionals over some subset of field configurations, chosen to vary in some suitably chosen way between points of the grid. (The details of how this is to be done don't

matter — i. e. have no consequences for the large-scale behaviour of the model — provided that specifying the value of a field configuration at all grid points is enough to identify it uniquely).

Such a theory is a close relative of the QM description of a crystal: in the latter, we specify the position and momentum of each crystal atom; in the former, the field strength and conjugate momentum at each gridpoint. (It is also very similar to the "lattice gauge theories" studied in lattice QCD.) It is also perfectly well-defined mathematically, and as such a valid entity to be viewed as approximately isomorphic to some subtheory of X.

We are interested in the structure of states of such theories, when studied at length-scales $\gg l$. To give a precise meaning to this, recall that the physical meaning of a state is given by its expectation value on elements of the algebra of observables. We can define the *large-scale* observables of a discretised QFT as those which are averages of grid observables over regions which are very large compared to the grid spacing; then we can regard the large-scale structure of a state as specified by its expectation value with respect to all large-scale observables.[6] The small-scale structure of a state is not to be treated as physically significant, since it is defined at lengthscales at which we expect the approximate isomorphism between X and our discretised QFT to break down. (Analogously, the structure of a classical phase-space distribution can be fairly reliably viewed as telling us about the actual world if it is studied on action scales large compared to $\hbar$, but is probably meaningless on scales small compared to $\hbar$).

Now, obviously there is a very large number of possible discretisations of a field theory: one for each possible choice of grid, at least (and no claim is being made here that the method above is the only way of implementing the cutoff). But renormalisation tames this profusion of theories to some extent, for it implies that the large-scale structure of two different discretisations can be made to be virtually identical (and totally identical in the limit as the two theories are studied on larger and larger length-scales) by finite rescaling of the field observables and finite adjustments in finitely many parameters of the Hamiltonians of the theories.

By analogy with section 2.2, we can construct an algebra net by associating to each region $\mathcal{O}$ which is large compared with the cutoff length l, an algebra $\mathcal{A}(\mathcal{O})$ of large-scale observables localised in $\mathcal{O}$. The statement that two discretised QFTs are structurally equivalent at large scales is then equivalent to the statement that both generate the same large-scale nets — a large-scale variant of the Net Assumption of section 2.2.

[6] By 'grid observables' I mean the basic field observables $\widehat{\phi}(x)$ and $\widehat{\pi}(x)$, not their pointwise products — so for instance $\int \mathrm{d}x\,\mathrm{d}y\, f(x)g(y)\widehat{\phi}(x)\widehat{\phi}(y)$ is a large-scale observable (if f and g vary slowly compared to l) but $\int \mathrm{d}x\, f(x)\widehat{\phi}(x)^2$ is not, irrespective of how smoothly f varies. This means, in particular, that the energy density is not a large-scale observable. In fact, the energy density of a discretised QFT is (a) cutoff-dependent, and (b) extremely large: this is the origin of the 'cosmological constant problem'.

In fact, we can mimic the algebraic description of section 2.2 further, by defining our cut-off QFTs directly in terms of their algebraic structures: we define a *scale-l algebraic QFT* as a map from regions of $\mathcal{M}$ which are large compared with l, to operator algebras, such that the axioms of AQFT are approximately satisfied on scales large compared with l. (Obviously this is not intended to be precise.) Two 'concrete' QFTs, such as two versions of the discretised QFT above, are then *scale-l equivalent* if they generate the same scale-l AQFT; hence scale-l equivalence is a generalisation of Borchers equivalence.

The results of renormalisation theory can now be simply stated: two discretisations of the same renormalisable QFT can be made scale-l equivalent just by finite renormalisations of the coefficients in one of them. So, for ϕ^4 QFT (for instance) there is a two-parameter family of scale-l AQFTs, parametrised by the renormalised mass and coupling constant.

Now if we know that a given scale-l AQFT is empirically accurate, we can find any number of concrete QFTs which are members of its scale-l equivalence class. We can choose any such member to calculate with, and make our choice based upon calculational convenience, but we need not regard any of them as candidates for the fundamental theory of nature. Rather, they are to be viewed as candidate theories to be approximately isomorphic to a subtheory of X.

From this perspective, the profusion of discretisations is no embarassment. For suppose QFT_1 and QFT_2 are members of the same scale-l equivalence class, and suppose the scale-l AQFT corresponding to that equivalence class accurately describes our observations (on scales large compared with l). Then it follows that QFT_1 and QFT_2 are both approximately isomorphic to a subtheory of X[7] with that approximate isomorphism breaking down around the cutoff lengthscale. But if so, then there can be no fact of the matter as to which of QFT_1 and QFT_2 is 'correct': *both* are approximately instantiated by X. If X is itself a QFT (presumably some interacting QFT which exactly satisfies the axioms of AQFT) then X itself will be inside the scale-l equivalence class, but if X is something else — a string theory, or some theory in which spacetime is quantized — then no elements of the equivalence class are fundamental: all are useful purely because of their structural resemblance to a subtheory of X.

3.4 Covariance

One of the least aesthetic features of discretised QFTs is their non-covariance, which manifests itself in (at least) two ways. Firstly, a lattice is inevitably not preserved under Poincare transformations (indeed, it is not even preserved under translations). Secondly, the commutators of field operators which are spacelike separated are in general not zero

[7] Note that QFT_1 and QFT_2 must be approximately isomorphic to the *same* subtheory of X, since they are intended to describe the same physical domains as one another.

(although they drop off extremely rapidly if the field operators are separated by distances large compared with the cutoff distance). This second effect is because of the discretisation of the Hamiltonian, in which spatial derivatives like $\nabla\phi(\mathbf{x})$ are replaced with terms like $(\phi(\mathbf{x}) - \phi(\mathbf{x}+\mathbf{l}))/l$, where l is the cutoff length and $\mathbf{l}$ is a vector of length l. As a consequence, influences can propagate instantaneously from $\mathbf{x}$ to $\mathbf{x}+\mathbf{l}$, although the effect is miniscule for widely separated points.

This is not to say that the approach advocated in section 3.3 implies that the actual world has faster-than-light signalling or a preferred reference frame. Recall that in this approach, our discretised QFT is taken to describe the actual world only insofar as it is isomorphic to a subtheory of the deeper "theory X", and we expect that this isomorphism will hold only for expectation values of field observables averaged over regions large compared with the cutoff length. But for these expectation values, we find that violations of covariance are very small, tending to zero as the regions over which the observables are averaged increase in size. In other words, non-covariance is essentially a small-scale property of discretised QFTs, and such properties are precisely those for which, it was argued, we should not regard QFT as telling us about the actual world.

Notwithstanding this, the situation is still rather unsatisfactory, at least if we view covariance as a fundamental property of physics rather than some effective limit of a non-covariant theory. There are a number of ways of coming to terms with this:

1. We might manage to find a covariant way of defining discretised QFTs. Obviously this would remove the covariance problem; however, it seems rather unlikely since cutoffs are about short lengthscales, and length is not a relativistic invariant. (Of course, if we look for a covariant QFT defined without a cutoff, we are moving away from the approach described here, and back towards the AQFT program.)

2. We might bite the bullet and accept that covariance is only an approximation. This would follow, for instance, if X were itself a non-covariant theory (such as Barbour (1994a, 1994b, 1999) has advocated in quantum gravity).

3. We might find that X is itself a totally covariant theory. There seems no reason why such a theory should not have a family of useful non- covariant approximations.

4. Perhaps most interestingly, X may be a theory which does not involve spacetime at all at a fundamental level (as is conjectured would be the case in loop-space quantum gravity, or non-perturbative string theory). In this case, relativistically covariant spacetimes would have to emerge as effective, approximate concepts from such a theory, and it would seem to be a curiosity rather than a pathology if this emergence proceded via non-covariant intermediate theories. After all, when we say that we expect spacetime to be quantized, we generally do not mean that it is literally to be divided into grid squares, but rather that it is to be replaced

in toto with some fundamentally discrete entity which in some circumstances is approximately isomorphic to spacetime — yet in recovering that spacetime it might be necessary to use an intermediate step which does involve treating spacetime as a grid. (To take an analogy, we think that the phase space of classical mechanics breaks down at scales of $\sim \hbar$, but that doesn't mean that we think that phase space really is a grid of $\hbar$-sized squares.)

4 Inequivalent representations

In section 2.2, we defined algebraic QFTs in terms of a net of *operator* algebras. The usual definition, however, is of a net of *abstract* algebras, which can be represented as Hilbert space operators in a number of inequivalent ways. At first sight these inequivalent representations seem to cause problems for QFT since there seems no principled way to select the 'right' one; it is also not immediately obvious what has happened to the inequivalent representations in our previous quantization (in section 2.3) of QFT in terms of wave-functionals on configuration space. In this section we shall address these issues, and will show that inequivalent representations — although important and interesting in QFT — pose no problems for its foundations.

4.1 The algebraic approach

Consider the quantum mechanics of some finite number of nonrelativistic, scalar particles. We could specify this theory by starting with the Hilbert space, and then describing the actions of the various observables (position, momentum and functions thereof) on the space. We would need to specify a specific function of the observables as the Hamiltonian, so as to tell us how states evolve in time.

Once we have these observables, we could shift to the Heisenberg representation by using the Hamiltonian to time-evolve the observables:

$$\widehat{X}(t) = \exp(-it\widehat{H})\widehat{X}\exp(+it\widehat{H}).$$

(The Hamiltonian would then be unnecessary, since all the information about the system's time evolution would be contained within the time-evolved observables.) Let $\mathcal{A}$ be the algebra of these observables.[8]

Now suppose we take $\mathcal{A}$ considered as an abstract algebra (with the norm topology, i. e. a C^*-algebra) and discard the Hilbert space. We would still have a great deal of dynamical information encoded in the algebraic and topological structure of the algebra, but would have lost the action of the algebra on physical states.

[8] If we were being mathematically rigorous, we would instead work with the algebra of *bounded functions* of the observables, since $\widehat{X}$ and $\widehat{P}$ are unbounded and thus very awkward to work with.

Suppose we try to recover the Hilbert space action. That is, we wish to find an isomorphism between the abstract algebra $\mathcal{A}$ and some subalgebra of the bounded operators on a Hilbert space $\mathcal{H}$, i. e. a *representation*[9] of $\mathcal{A}$. If this representation turns out to be unique up to unitary equivalence (an assumption which was often made uncritically in the early days of quantum theory) then we must have recovered the original Hilbert space, i. e. the representation must be isomorphic to the original choice of Hilbert space and of observables on it.

In fact, for a system of finitely many particles the representation is indeed unique — provided we require it to be *irreducible*. This last is the requirement that the Hilbert space does not split into orthogonal subspaces preserved by the action of $\mathcal{A}$. Equivalently, there must exist in the Hilbert space a *cyclic vector* whose image under the action of $\mathcal{A}$ is dense in the Hilbert space. Physically this requirement rules out the existence of superselection rules: superselection sectors can be identified with the irreducible sectors of a reducible representation.

By an *algebraic* quantum theory, we shall mean one which is specified by giving the algebraic structure of the observables but not the Hilbert space on which they act. A theorem of Stone and Von Neumann tells us that any such quantum theory with a *finite* number of degrees of freedom — i. e. finitely many 'position' and 'momentum' observables, corresponding to a classical theory with a finite-dimensional phase space — has a unique irreducible representation (up to unitary equivalence). Hence for such systems the algebraic and Hilbert-space ways of specification are equivalent.

4.2 Infinite-dimensional systems

When the number of degrees of freedom of a theory becomes infinite, the Stone-von Neumann theorem fails. Such theories generally have a large number of inequivalent representations, and so giving the algebraic structure does not completely specify the theory.

Nonetheless it still seems natural to specify a theory algebraically, because

1. The way in which we usually describe a system is to describe its observables. Furthermore, a theorem of Fell (see below) tells us that no amount of empirical data can tell us what the 'real' representation is.

2. As explained in section 2.2, the 'quantization' process begins with the classical observables and tries to map them onto quantum-mechanical operators. Since in this process we are trying to impose purely algebraic restrictions on this map (specifically, the restriction that Poisson brackets go over to commutators) it will generally specify

[9]Technically, a representation need not be an isomorphism. However, here we are working with simple C^*-algebras, for which all irreducible representations are faithful (i. e. isomorphic to the original algebra).

> only the algebraic structure of a quantum theory, with an ambiguity left as to the Hilbert space action. In other words, two quantum theories which are algebraically identical could be regarded as equally valid quantizations of the classical system.[10]

If we were to take a robustly instrumentalist viewpoint there would be no particular problem: Fell's theorem says that a finite number of measurements, each conducted with finite accuracy, cannot distinguish between representations. Hence we could reproduce any experimental results using whichever representation was most convenient.[11]

From a realist standpoint, however, we appear to have a dilemma. Even if no empirical data lets us distinguish between representations, nonetheless there should be a fact of the matter as to which is the 'correct' representation; yet it is provably impossible for us ever to discover this fact. To see this, suppose we have two representations on Hilbert spaces $\mathcal{H}_1$ and $\mathcal{H}_2$, and suppose the 'real' representation is the first. Then if the system is in some state $|A\rangle$, Fell's theorem tells us that

> For any operation $\mathcal{O}$ carried out with finite accuracy on $|A\rangle$, there is a state $|\mathcal{O};\mathcal{A}\rangle$ in $\mathcal{H}_2$ such that all data resulting from the operation are consistent with the real representation being on $\mathcal{H}_2$ and the real state being $|\mathcal{O};\mathcal{A}\rangle$.

In fact, this dilemma will prove only apparent, as we will find in the next three sections when we study representation ambiguities in more detail. To do so, we distinguish two ways in which inequivalent representations can occur: one associated with the short-distance and high-energy ('ultra-violet') degrees of freedom, one with the long-distance ('infra-red') ones. Given our tolerant attitude to short-distance cutoffs (c. f. section 3), it should come as no surprise that all the important sources of inequivalent representations of a given QFT fall into the second category: for the Stone-von Neumann theorem guarantees uniqueness of representation for any theory with finitely many degrees of freedom, and a field theory in any finite region has only finitely many degrees of freedom below a finite energy bound. We will argue that the ultra-violet representation ambiguity is a pure mathematical artefact, whilst the infra-red ambiguity is physically real but non-pathological.

4.3 Ultra-violet degrees of freedom

It is probably already clear to the reader why we can ignore the ultra-violet representation ambiguity: UV-inequivalent representations occur because of the existence of degrees of

[10] Separately from this, of course, we might not in general expect quantization even to produce a unique *algebraic* structure for the quantum observables, because of ambiguities about operator ordering.

[11] It is important in this approach that we have all the observables in the algebra. If, for instance, the (renormalized) stress-energy tensor is not in the algebra, then its values may distinguish representations (Wald 1994).

freedom at arbitrarily short lengthscales, and in section 3 it was argued that real QFTs are cut off at short, though finite lengthscales. In particular, the discretised QFTs of section 3.3 have only finitely many degrees of freedom per space-time point, and hence no UV-inequivalent representations. Since renormalisation theory tells us that any QFT is scale-l equivalent to some such discretised QFT (on lengthscales long compared to the grid size for the latter theory), it follows that any occurence of UV-inequivalent representations is purely a mathematical artefact and has no physical significance.

4.4 Inequivalent representations in non-relativistic physics

The physical significance of IR-inequivalent representations has long been appreciated in AQFT; it will be reviewed here from the perspective which this paper adopts towards QFT, but with no claim to originality. We begin by returning to non-relativistic QM: imagine a line of n two-state systems — spin-half particles fixed in place, say. The Hilbert space of each system is two-dimensional, and we can construct the overall Hilbert space by taking the n-fold tensor product. The resulting space will have dimension 2^n, which is as we expect since we must make n yes-no choices to choose a state with (say) given spin in the z-direction for each component system.

Now suppose our system becomes infinitely large. The dimension of the system will be infinite, of course, but it will be a larger infinity than those to which we are used — specifically it will be $2^{\aleph_0}$, the cardinality of the continuum, which is strictly larger than the cardinality $\aleph_0$ of the integers.[12]

It follows that systems with infinitely many components have a Hilbert space which is non-separable (i. e., has uncountable dimension). To see the consequences of this, consider the operator algebra of our set of two-state systems. It consists of the set of linear combinations of spin operators, and hence has countably many linearly independent elements. The action of this algebra on any given state will generate only countably many linearly independent states, hence the action of the operator algebra on the total, non-separable space is highly reducible.

In other words, the non-separable space has decomposed into various superselection sectors (uncountably many, in fact), each giving an irreducible representation of the operator algebra.

To see what these sectors are, suppose we start with all components having spin up. Then the action of any element of the algebra can, at most, cause finitely many components to have spin down. So no amount of algebraic action can transform such a state into one in which, say, every second component has spin up. This state, in turn, can be transformed into other states differing from it in finitely many places, but not into a state in which all components are spin down . . . or every third component is spin down

[12] See, e. g., Enderton (1977).

... or where half the states are spin up but the spin-up states are grouped in pairs ...

In other words, the different representations describe states which are 'infinitely' different from one another (see Morrison (1962) for a full discussion). In statistical mechanics, different representations describe different phases, since at arbitrarily large length-scales — tending to the thermodynamic limit — only the infinite differences between representations remain visible.

For such systems, of course, the use of different representations is an idealization, since the actual systems are finite.

4.5 IR-inequivalence in quantum field theory

Guided by the nonrelativistic example, we look for inequivalent representations in field theory by looking for states differing on asymptotically large scales. The difference from the previous section is that here infinite systems are perfectly respectable: in fact, our discretised QFTs are effectively infinite crystals.

This resolves the problem mentioned at the start of section 4: what is the connection between inequivalent representations and the wave-functional approach of section 2.3? The answer is that the space over which the wave-functionals are defined — the field-configuration-space $\mathcal{S}$ — is infinite-dimensional both because functions may vary on arbitrarily short length-scales, and because they may have arbitrary large-distance boundary conditions. The former reason is nullified when a short-distance cutoff is imposed (as in section 3) but the latter one also interferes with the definition of the functional integral, and has to be dealt with by imposing boundary conditions at infinity (such as the requirement that functions in $\mathcal{S}$ be square- integrable). Each choice of boundary condition generates a different, and inequivalent, representation.

Examples of inequivalent representations are then:

Differing mass-densities at infinity

The closest analogue to the previous section is the case of states with asymptotically nonvanishing mass density: average density is superselected (Haag 1996) and different densities correspond to different representations. It is worth noting that this rules out the Fock representation for an open universe, even a flat one, since any realistic state of such a universe will have nonzero mass density; the Fock representation, which has finite-energy states, stands in the same relation to a realistic open-universe representation as an asymptotically empty classical spacetime does to a realistic open cosmology (Wald 1984).

The overall velocity of the universe is also superselected, of course, except in the special case where the (expected value of the) stress-energy tensor is asymptotically diagonal.

Differing total charge

Even when the spacetime is asymptotically empty, inequivalent representations exist. The total charge of a system leaves its imprint on the spacetime at arbitrarily long distances in the form of the asymptotic Coulomb field, whose flux through a sphere is independent of the size of that sphere. As such, total charge is superselected (Strocchi and Wightman 1976). (Another way to see this is to remember that field lines extend to infinity, so no amount of local operations will remove them.)

In fact, the centre-of-mass velocity of a charged system is also superselected (Haag 1996), since the asymptotic field will be a Lorentz-boosted Coulomb field and the associated magnetic field lines also extend to infinity.

Inequivalent vacua

A major component of the Standard Model of particle physics (Peskin and Schroeder 1995) is the presence of field theories with degenerate vacua. Classically this corresponds to the existence of nonzero field configurations which are global minima of the energy. If the field theory has a continuous global symmetry then there will exist a continuum of vacuum states, wave-packets around which will give us the quantum ground states. Since these states differ from one another everywhere in space, they too belong to inequivalent representations.

Changes in the interaction parameters

A famous theorem (Haag's theorem; see Haag 1996) tells us, in effect, that interacting fields cannot be represented on the same Hilbert space as the corresponding free fields, even when we consider asymptotically early or late field operators. In the current framework this may be understood as follows: introducing the interaction will change the ground state everywhere in space, causing it to be infinitely different from the free-field ground state and so in a different superselection sector.

In fact, and as these examples show, there is not really anything 'quantum' about IR-inequivalent representations: in many cases, picking a representation is rather like picking boundary conditions at infinity in a classical problem. The following classical — and utterly banal — version of Fell's theorem shows why the 'dilemma for realism' mentioned in section 4.1 is nothing of the kind as far as IR-inequivalent representations are concerned:

> **Classical Fell 'theorem':** the spatial average of a classical quantity over an infinite universe cannot be empirically determined by measurements confined to a finite spatial region.

4.6 Retrospective

From a realist perspective the sting of the representation ambiguity has largely been drawn. Locally, any representation ambiguity is artificial, caused by the presence of unphysical degrees of freedom beyond the high-energy limit of the theory's validity. Globally, there may indeed be representation ambiguities — depending on cosmology, and the topology of the universe — but the inaccessible information which they encode is 'respectable', analogous to the classical inaccessibility of the long-distance structure of the universe.

Nonetheless it would be more aesthetic to lose the long-distance ambiguity as well. This can only be done if the universe is finite. There has long been a divide between observational cosmology (favouring an open universe) and theory (preferring closed universes); this analysis of representations makes a modest contribution to keeping the divide open.

For the purposes of this paper, representational ambiguities need no longer concern us. In practice we are always analysing a theory in a finite region — and idealizing the system beyond that region in whatever manner is convenient — so different choices of representation should not affect our conclusions.

5 Localisation in quantum field theory

The phenomenology of particle physics — and indeed, of virtually all of science — makes extensive use of the concept of localisation: that is, of the concept that physical systems have at least some states which localised in finite spatial regions. There are a number of results in AQFT which apparently rule out the possibility of such states, and this is sometimes described as a paradox in QFT. In this section the problem of localisation will be analysed from the viewpoint of the previous sections, and it will be argued that nothing paradoxical is going on.

In a relativistic theory, there are two natural ways of thinking about localisation: *spatial* localisation, where something is localised in some subregion Σ_i of a spacelike slice Σ (and therefore *not* localised anywhere else in Σ), and *spacetime* localisation, where something is localised in some subregion $\mathcal{O}$ of spacetime (and therefore *not* localised in any region of spacetime spacelike separated from $\mathcal{O}$). The latter concept makes sense in AQFT (including the scale-l approximate AQFTs of section 3.3); the former needs some concrete QFT to be interpreted. In this section we will move freely between the two notions, using a discretised QFT where necessary to make sense of spatial localisation, and relying on renormalisation theory to ensure that we are not making cutoff-dependent statements.

5.1 Localisation and the Reeh-Schlieder theorem

We have already established (in section 2.2) that in QFT the idea of localisation arises through the spatial localisation of the field observables $\widehat{\phi}(\mathbf{x})$ and $\widehat{\pi}(\mathbf{x})$ (here, for convenience, we work with a concrete element of the scale-l equivalence class of the QFT in question). But how are we to go from localised observables to localised states? We might begin by trying the following (phrased in terms of spacetime localisation):

> **Naive localisation**: A state $|\psi\rangle$ is localised in a region $\mathcal{O}$ iff $\langle\psi|\,\widehat{A}\,|\psi\rangle = 0$ for any observable $\widehat{A}$ localised in a region spacelike separated from $\mathcal{O}$.

This seems plausible when we compare to the classical case: there a state is localised in $\mathcal{O}$ if $\pi(x) = \phi(x) = 0$ for any x spacelike separated from $\mathcal{O}$. But it is mathematically impossible for any states to satisfy it, for it implies that for any such $x = (\mathbf{x}, t)$, and for any n,

$$\langle\psi|\,\widehat{\phi}^{\,n}(\mathbf{x}, t)\,|\psi\rangle = \langle\psi|\,\widehat{\pi}^{n}(\mathbf{x}, t)\,|\psi\rangle = 0. \tag{17}$$

But this would imply that $|\phi\rangle$ was a simultaneous eigenstate of $\widehat{\pi}(\mathbf{x})$ and $\widehat{\phi}(\mathbf{x})$, and these operators have no eigenstates in common (the mathematics, once we have introduced a cut-off to deal with operators defined at a point, is the same as for the nonrelativistic operators $\widehat{X}, \widehat{P}$, which are well-known to have no eigenstates in common).

Physically it is easy to see what is happening here. The vacuum state of a field theory (which we will denote by $|\Omega\rangle$) is not 'nothingness', or 'empty space'; it is simply a slightly colourful way of describing the ground state of the field's Hamiltonian. In solid-state systems this state is just the zero-temperature state of the solid, in which the atoms will not be at rest but will have zero-temperature fluctuations; the same will be true for the field excitations of a relativistic field theory.

This suggests, however, an alternative: rather than consider states which are local *simpliciter*, we could consider states which *differ* only locally. The vacuum could be used as a reference state, and states could be declared 'local' if they differed only locally from the vacuum.

To make this quantitative, we need to understand what it could mean for two quantum states to 'differ locally'. Recall that in section 2.3 we defined spatially local subsystems of the QFT, with Hilbert spaces $\mathcal{H}_{\Sigma_j}$, which described the degrees of freedom of the QFT localised in Σ_j; this formally defined concept can be exactly defined in discretised QFT, provided that the grid spacing is small compared to each Σ_j. If the vacuum were a product state with respect to any decomposition of the form $\mathcal{H}_\Sigma = \mathcal{H}_{\Sigma_1} \otimes \mathcal{H}_{\Sigma_2}$ (where $\Sigma = \Sigma_1 \cup \Sigma_2$ and $\Sigma_1 \cap \Sigma_2 = \emptyset$)[13] then defining 'differing locally' would be straightforward: the vacuum would have form $|\Omega\rangle = |\Omega_1\rangle \otimes |\Omega_2\rangle$, and any state $|\psi\rangle$ could be said to differ

[13] Since we are working with discretised QFT here, a more accurate statement would be 'all grid points in Σ are contained either in Σ_1 or Σ_2, and Σ_1 and Σ_2 do not overlap.'

from the vacuum only in Σ_1 if (a) $|\psi\rangle$ was also a product state, $|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle$, and (b) $|\psi_2\rangle = |\Omega_2\rangle$.

However, the vacuum is not a product state, but a highly entangled one (this can be readily seen by calculating expectation values such as

$$\langle\Omega|\,\widehat{\phi}(x)\widehat{\phi}(y)\,|\Omega\rangle - \langle\Omega|\,\widehat{\phi}(x)\,|\Omega\rangle\,\langle\Omega|\,\widehat{\phi}(y)\,|\Omega\rangle \tag{18}$$

and showing that they are non-zero for all spacelike separated x, y,[14] and is discussed by Clifton and Halvorson (2001).) As such, defining 'differing locally' is somewhat more subtle, since it is difficult and even controversial to say what the state of a subsystem of a quantum system is, when the total state of the system is entangled. Some possible definitions might be:

1. Two states might be said to differ only within Σ_1 if their expectation values coincide with respect to all operators localised outside Σ_1.

2. We might take the 'state of a subsystem' to be the density operator obtained when all degrees of freedom of the other subsystems are traced over. In this case, two states $|\psi_a\rangle$ and $|\psi_b\rangle$ could be said to differ only within Σ_1 if (with $\mathsf{Tr}_{\mathcal{H}_{\Sigma_1}}$ being tracing over degrees of freedom inside Σ_1)

$$\mathsf{Tr}_{\mathcal{H}_{\Sigma_1}}\,|\psi_a\rangle\,\langle\psi_a| = \mathsf{Tr}_{\mathcal{H}_{\Sigma_1}}\,|\psi_b\rangle\,\langle\psi_b|\,. \tag{19}$$

3. The problem with the above definition of the state of a subsystem is that we cannot then recover the state of the whole system from the states of its components. This could be taken to indicate the intrinsic nonlocality of quantum states; this view was recently rejected by Deutsch and Hayden (2000), who propose that locality in quantum mechanics is best understood in the Heisenberg picture. They pick a specific Hilbert-space vector as a reference state, to be used to calculate all expectation values irrespective of the particular initial conditions and then treat the quantum state as being specified at time t by the set of all Heisenberg operators at time t.

[14] For a free scalar field, with $x = (\mathbf{x}, t)$ and $y = (\mathbf{y}, t)$ we get

$$\langle\Omega|\,\widehat{\phi}(x)\widehat{\phi}(y)\,|\Omega\rangle - \langle\Omega|\,\widehat{\phi}(x)\,|\Omega\rangle\,\langle\Omega|\,\widehat{\phi}(y)\,|\Omega\rangle = \frac{1}{2}(m^2 - \nabla^2)^{-\frac{1}{2}}\delta(\mathbf{x} - \mathbf{y});$$

smearing $\widehat{\phi}(\mathbf{x}, t)$ and $\widehat{\phi}(\mathbf{x}, t)$ out with non-overlapping spatial test-functions f, g, this becomes $\frac{1}{2}\int_\Sigma f(\mathbf{x})[(m^2 - \nabla^2)^{-\frac{1}{2}}g](\mathbf{x})$. That this is nonzero for generic f, g, irrespective of the spatial separation of their supports, follows from the antilocality of the operator $(m^2 - \nabla^2)^{-\frac{1}{2}}$ (Segal 1964). Adding interactions to the free field will lead to perturbative modifications to these results but should not change their qualitative nature, provided the interactions are weak enough at large distance-scales to treat perturbatively — hence adding a $\lambda\phi^4$ interaction will have only quantitative effect, but adding a coupling to a non-Abelian gauge field may have more drastic consequences.

The obvious definition of 'differing locally' within this framework is that two states differ only within Σ_1 if they have the same Heisenberg operators outside Σ_1.

4. Leaving aside the question of how to specify the state of a subsystem, we might define two states as differing only within Σ_1 if they are connected by a unitary operator localised within Σ_1 (such an operator will have form $\widehat{U}_{\Sigma_1} \otimes \widehat{\mathbf{1}}_{\Sigma_2}$ acting on the tensor-product space $\mathcal{H}_\Sigma = \mathcal{H}_{\Sigma_1} \otimes \mathcal{H}_{\Sigma_2}$).

Fortunately, all of these definitions coincide[15] and this in turn supports the naturalness of each of them as a definition of 'differing locally'. Using the vacuum as our reference state, and using for convenience the first definition, we obtain our definition of localised states, which we refer to as Knight localisation after Knight (1961) who first proposed such a definition. We state it in the more general context of AQFT (for which reason we return to spacetime, as opposed to spatial, localisation):

> **Knight localisation:** a state $|\psi\rangle$ is localised in a region $\mathcal{O}$ iff $\langle\psi|\,\widehat{A}\,|\psi\rangle - \langle\Omega|\,\widehat{A}\,|\Omega\rangle = 0$ for any observable $\widehat{A}$ localised in a region spacelike separated from $\mathcal{O}$.

However, Knight localisation differs in one important respect from the sort of localisation which we encounter in NRQM. In the latter, properties like 'is localised in $\mathcal{O}$' are treatable in the same way as properties like 'has energy E' or 'has momentum less than p': that is, we can define a projection operator whose intended interpretation is 'localised in $\mathcal{O}$', whose range is the space of all such states. This would be possible for Knight-localised states iff they form a subspace: that is, iff any superposition of two states Knight-localised in $\mathcal{O}$ is also Knight-localised in $\mathcal{O}$.

The fact that Knight-localised states do not have this property is a consequence of the Reeh-Schlieder theorem (Reeh and Schlieder 1961).

> **Reeh-Schlieder theorem**: for any region $\mathcal{O}$, the set of vectors $\mathcal{A}(\mathcal{O})\,|\Omega\rangle$ generated by the action of operators localised within $\mathcal{O}$ upon the vacuum, spans the Hilbert space of the QFT.

(For a proof, and further discussion, see Haag 1996.) It follows[16] from the Reeh-Schlieder theorem that states Knight-localised at $\mathcal{O}$ span the entire state space, which rules out any possibility of a projector meaning 'localised with certainty in $\mathcal{O}$'.

[15] The equivalence of (1), (2) and (4) is trivial; the equivalence of (3) and (4) is given in Deutsch and Hayden 2000 (and is, again, trivial once Deutsch and Hayden's formalism is understood).

[16] To see that it follows, we need only note that the unitary elements of a (bounded) operator algebra $\mathcal{A}(\mathcal{O})$ span $\mathcal{A}(\mathcal{O})$. This can be proved as follows: for any bounded Hermitian element $\widehat{H}$ of $\mathcal{A}(\mathcal{O})$, and any $t \neq 0$, $(it)^{-1}(\exp(it\widehat{H}) - \widehat{1})$ is a linear combination of unitary elements of $\mathcal{A}(\mathcal{O})$. As $t \to 0$, this sequence tends to $\widehat{H}$, hence $\widehat{H}$ is in the span of the unitary operators. To complete the proof, simply recall that any linear operator can be written as $\widehat{A} + i\widehat{B}$, where $\widehat{A}$ and $\widehat{B}$ are Hermitian.

It has long been understood that the Reeh-Schlieder theorem is a consequence of the entanglement of the QFT vacuum. Non-relativistic quantum mechanics furnishes us with many examples where, given an entangled state of the Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, unitary operations on $\mathcal{H}_A$ alone suffice to produce a basis for $\mathcal{H}_A \otimes \mathcal{H}_B$. One example, which plays an important role in quantum teleportation (Bennett, Brassard, Jozsa, Peres, and Wootters 1993), is the *Bell basis*: we take a system of two qubits (i. e., two-state systems) and prepare them in one of these four states

$$\begin{array}{rcl} |B1\rangle & = & \frac{1}{\sqrt{2}}\left(|0\rangle\otimes|0\rangle + |1\rangle\otimes|1\rangle\right) \\ |B2\rangle & = & \frac{1}{\sqrt{2}}\left(|0\rangle\otimes|0\rangle - |1\rangle\otimes|1\rangle\right) \\ |B3\rangle & = & \frac{1}{\sqrt{2}}\left(|1\rangle\otimes|0\rangle + |0\rangle\otimes|1\rangle\right) \\ |B4\rangle & = & \frac{1}{\sqrt{2}}\left(|1\rangle\otimes|0\rangle - |0\rangle\otimes|1\rangle\right) \end{array}. \tag{20}$$

Obviously this is a basis for the combined two-qubit system, but also any element of it is cyclic under unitary operations carried out on the first qubit alone, as is easily shown:

$$\begin{array}{rcl} (M_A \otimes \widehat{1})\,|B1\rangle & = & |B2\rangle \\ (M_B \otimes \widehat{1})\,|B1\rangle & = & |B3\rangle \\ (M_C \otimes \widehat{1})\,|B1\rangle & = & |B4\rangle \end{array} \tag{21}$$

where

$$M_A = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}; M_B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}; M_C = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

in the $(|0\rangle\,, |1\rangle)$ basis. In other words, even if the qubits are macroscopically separated, if they begin in an appropriately entangled state then it is possible by operating on one qubit to generate a set of states which span the Hilbert space. (Examples of this kind are analysed in rather more detail by Redhead (1995) and Clifton and Halvorson (2001).)

Since (as discussed above) the QFT vacuum is entangled, the Reeh-Schlieder theorem should come as no surprise to us. Nor does it cause any *logical* problems for Knight's definition of localisation: there are many perfectly respectable properties of quantum states which are not preserved under linear superposition, such as being an eigenstate of energy, or being an entangled state.

It does, however, cause us *practical* problems. The sort of localisation which we use in NRQM *is* preserved under linear superpositions, and this fact is essential to the analysis of NRQM problems — so, at least on the face of it, NRQM localisation and Knight localisation must be different concepts. Nor is the problem confined to NRQM: in scattering theory, for instance, it is crucial to be able to discuss the amplitude for a particle to be scattered to a given area, and such discussions presuppose that spatial localisation is preserved under superpositions.

So what is going on? In the face of the Reeh-Schlieder theorem, there is little prospect of finding an alternative to Knight localisation which does have the required properties (see Halvorson and Clifton (2001) for a wide variety of no-go theorems); so, how are NRQM, and scattering phenomenology, compatible with QFT?

The key to this question is the fact that NRQM is not supposed to be *perfectly* compatible with QFT: rather, it is supposed to be an approximation to QFT valid only in certain regimes of QFT. (And the same is true for scattering theory, which is a marvellous tool to analyse high-energy collisions, but of limited use in understanding, e. g. , quark confinement.) With this in mind, we will lower our sights and seek an approximate form of locality which, in certain domains of QFT, will be approximately preserved under linear superposition. Section 5.2 will construct such an approximation, and section 5.3 will address the question of whether it is after all adequate for our needs.

5.2 Effective localisation

Our approximate concept of localisation will be defined as follows (it is again defined in terms of spatial localisation, partly for ease of comparison with nonrelativistic localisation):

1. Effective localisation (qualitative form): A state $|\psi\rangle$ is *effectively localised* in a spatial region Σ_i iff for any function $\widehat{f}$ of field operators $\widehat{\phi}, \widehat{\pi}$, $\langle\psi|\,\widehat{f}\,|\psi\rangle - \langle\Omega|\,\widehat{f}\,|\Omega\rangle$ is negligibly small when $\widehat{f}$ is evaluated for field operators outside Σ_i, compared to its values when evaluated for field operators within Σ_i.

2. The effective localisation principle (ELP) (qualitative form: A subspace $\mathcal{H}$ of the QFT Hilbert space $\mathcal{H}_\Sigma$ obeys the ELP on scale L iff for any spatial region $\mathcal{S}$ large compared with L, a superposition of states effectively localised in $\mathcal{S}$ is effectively localised in effectively the same region.

These qualitative notions can be made precise in a number of ways, such as:

1. Effective Localisation (quantitative form): A state is *L-localised* in a region Σ_i, iff for any function $\widehat{f}$ of field operators $\widehat{\phi}, \widehat{\pi}$, $\langle\psi|\,\widehat{f}\,|\psi\rangle - \langle\Omega|\,\widehat{f}\,|\Omega\rangle$ falls off for large d like (or faster than) $\exp(-d/L)$, where d is the distance from Σ_i at which the function $\widehat{(f)}$ is evaluated. (Note that there is no difference, according to this definition, between a state L-localised at some spatial point $\mathbf{x}$ and a state L-localised in a region of size $\sim L$ around $\mathbf{x}$.)

2. ELP (quantitative form): A state obeys the ELP on scale L iff, for any 3-sphere $\mathcal{S}$ of radius $> L$, a superposition of states L-localised in $\mathcal{S}$ is L-localised in $\mathcal{S}$.

Obviously ELP cannot hold on any scale for the Hilbert space $\mathcal{H}_\Sigma$ of the full QFT, on pain of violating the Reeh-Schlieder theorem (the space of states effectively localised in a region obviously includes all those Knight-localised in a region, and as mentioned in section 5.1, the span of all such states is the whole of $\mathcal{H}_\Sigma$.) However, if some subspace $\mathcal{H}$ of $\mathcal{H}_\Sigma$ is such that:

1. ELP holds, for all regions large compared with some lengthscale L;
2. the current state of the QFT lies within $\mathcal{H}$; and
3. $\mathcal{H}$ is approximately conserved, on timescales of interest to us, by the QFT dynamics

then we will, approximately, be able to define projection operators in $\mathcal{H}$ which project onto states localised in a given region, provided that that region is large compared to L. (Note that L is nothing to do with the cutoff lengthscale l, except that L must be large relative to l for our results to be cutoff-independent.)

Why should we expect ELP to hold for any subspaces of $\mathcal{H}_\Sigma$? To answer this, we need to analyse the entanglement of the vacuum in somewhat more detail. For if it were completely non-entangled, ELP would indeed hold for all of $\mathcal{H}_\Sigma$, as argued for above. We might expect, then, that if the vacuum entanglement drops off rapidly above certain distance-scales, then above those scales localisation would be 'almost' preserved under superposition, and only extremely careful superposition of states would spoil this; hence, we would expect it to be fairly typical for a given subspace of $\mathcal{H}_\Sigma$ to satisfy ELP on these lengthscales.

Does vacuum entanglement in fact decrease with distance; equivalently, do the correlations in the vacuum decrease with distance? For a massive QFT it can be shown that they drop off exponentially, with the lengthscale given by the 'Compton wavelength' $1/m$ (where m is the renormalised mass term in the QFT Hamiltonian): this is provable rigorously for any QFT satisfying the AQFT axioms (Fredenhagen 1985), and can be calculated perturbatively in the case of QFTs with weak large-scale interaction terms. (QCD, and other theories with strong long-distance interactions, may not be covered by either case — rather little is known about the QCD vacuum, in fact.)

To see intuitively the significance of the $1/m$ lengthscale, recall that the Hamiltonian for a free scalar QFT has the form

$$\widehat{H}\Psi[\chi] = \frac{1}{2}\int_\Sigma \mathrm{d}^3\mathbf{x}\,\left(\frac{\delta^2}{\delta\chi(\mathbf{x})^2} + (\nabla\chi)(\mathbf{x})^2 + m^2\chi(\mathbf{x})^2\right)\Psi[\chi], \tag{22}$$

and that in a discretised QFT we replace derivative terms like $\nabla\widehat{\phi}(\mathbf{x})$ with terms like $(\phi(\mathbf{x}) - \phi(\mathbf{x}+\mathbf{l}))/l$, where $|\mathbf{l}| = l$ and l is the cutoff length. These discretised derivative terms are the only ones in the Hamiltonian which lead to entanglement, as they are the only ones which couple observables at different lattice points.

Now there is no requirement that $l \ll 1/m$ — all we require of l is that it is small compared with the lengthscales at which we want to study the QFT's structure. Admittedly, if l is smaller than, or comparable to, $1/m$, then the coupling terms $(\phi(\mathbf{x}) - \phi(\mathbf{x}+\mathbf{l}))/l$, which are responsible for entanglement, will be significant relative to the other terms in the Hamiltonian. However, as l becomes large compared with $1/m$, the coupling terms become only a small perturbation on the Hamiltonian

$$\widehat{H}\Psi[\chi] = \frac{1}{2}\int_\Sigma \mathrm{d}^3\mathbf{x}\,\left(\frac{\delta^2}{\delta\chi(\mathbf{x})^2} + m^2\chi(\mathbf{x})^2\right)\Psi[\chi], \tag{23}$$

which is the Hamiltonian of a set of uncoupled oscillators, one at each lattice point. Thus, on such scales we should expect to find the ground state only slightly entangled.

So, we expect that for a massive QFT with mass m, we should find it to be common for the ELP to hold for subspaces on lengthscales large relative to $1/m$. Obviously, this heuristic argument has to be tested for any specific subspace in which we are interested. In particular, we find it to hold for the one-particle sector of any free (or asymptotically free) massive bosonic QFT, suggesting that such theories have a quite satisfactory notion of particle localisation (this result is proved, and analysed *in extenso*, in Wallace (2001a)).

For a QFT which describes some massless fields, the vacuum correlation functions still drop off with distance, but follow a power-law dropoff rather than an exponential one; hence there is no characteristic lengthscale on which correlations exist. This does not rule out finding subspaces of the Hilbert space of the QFT for which ELP applies on some useful scale; however, it suggests that the scale will vary according to the specific problem being examined.

Furthermore, we expect there to be physically relevant sectors of at least some massless QFTs in which states are localised only on extremely large lengthscales (if at all). For instance, there should be a low-energy sector of QED in which electron number is approximately conserved, in which the electromagnetic field has effectively no independent degrees of freedom, and in which each electron has an associated Coulomb field. In such sectors, the state of the electromagnetic field is not really localised at all, or at least not on any characteristic lengthscale; analysis of such sectors would be necessary to establish the validity of ELP for nonrelativistic, interacting particles. (See chapter VI of Haag (1996) and references therein for further discussion of the state space of QED, from the perspective of AQFT.)

5.3 The conceptual justification of effective localisation

We have now constructed a definition of localisation which approximately has the properties we wish: that is, it designates a set of states which are approximately localised, and that set is closed under linear superpositions. But is this approximate localisation good enough? After all, in general 'effectively localised' states are not exactly localised — particle states, in particular, invariably differ from the vacuum everywhere in space. It follows that particle creation and annihilation operators cannot be exactly localised, which seems to preclude localised particle-detecting devices.

This problems has recently been addressed from an AQFT viewpoint by Halvorson and Clifton (2001) (see also Haag (1996) for a more technical discussion along similar lines). Their argument is as follows:

1. What we actually measure are (exactly) local operators.

2. Although particles cannot be detected by measurements of any local operator, they can be detected by measurements of operators which are very close (in operator norm) to local operators.

3. It follows that to a high degree of accuracy, we *can* detect particles; however, the detection will not be 100% reliable.

In their approach,

> It is not (strictly speaking) true that we observe particles. Rather, there are 'observation events', and these observation events are consistent (to a good degree of accuracy) with the supposition that they are brought about by (localizable) particles.(Halvorson and Clifton 2001)

From the perspective of this paper, the problem with this approach is its a priori assumption that what we measure are always exactly localised operators. This is, of course, an interpretive axiom of AQFT as it is often presented, but it effectively assumes the presence of outside observers whose measurements cannot be treated within the ordinary dynamics of the QFT. We shall instead construct an account which treats observers as part of the internal dynamics of the system (although, apart from that difference of emphasis, the solution below will be rather similar in character to that of Halvorson and Clifton).

If we wish instead to treat our QFT as a closed system, and the measurement process as part of the internal dynamics of that system, then it is an open question whether or not we must treat question our measurements as localised. Furthermore, we have at least some reason to think that the answer to the question is negative — for we believe our measuring devices (including ourselves) to be made out of particles, and we have already noted the fact that particles are never exactly localised.

But does it even make sense to consider non-localized measuring devices? To see that it does (at least in some sense), we return (again) to non-relativistic quantum mechanics. In NRQM it is unproblematic, at least in principle, to construct measurement devices out of atoms. Such devices usually have rather large masses, so despite the uncertainty principle they can be put in states whose position and momentum are both very sharp.

In particular, if a measurement device's momentum is sharply peaked around zero then the system will effectively remain in a single well-determined region.

But will it exactly remain in that region? The mathematics says not, for it is mathematically impossible for any centre-of-mass wavefunction to remain localised in a finite spatial region for any finite period of time. (A rather general proof of this result has been given by Hegerfeldt (1998a, 1998b), and is discussed by Halvorson and Clifton (2001).) This is implicit, in fact, in the wave-functions we usually use in NRQM to describe 'localised' systems: they are generally described by fairly sharp Gaussians, but any Gaussian — no matter how sharp — is non-zero everywhere in space.

Does this mean that the Geiger counter in my lab has finite amplitude to be in Andromeda? No, for the representation of a quantum object by a Gaussian state relies on the fact that the object is not entangled with its environment. But the part of the wave-function representing the device as (for instance) in the walls of the lab will obviously interact with its environment in a rather different way from the part which represents the device as in the middle of the lab. As such, the device should not strictly speaking be represented by a pure state at all, but is instead entangled with its surroundings.

The reason why we do not in practice need to allow for this entanglement is that it is ludicrously small. Consider, for instance, a device of mass m whose centre-of-mass wave-function has Gaussian form

$$\psi \propto \exp(-r^2/2L^2). \tag{24}$$

If the device is allowed to evolve freely, after some time t the centre-of-mass wave-function will evolve to

$$\psi(r,t) \propto \exp(-r^2/2L_t^2) \tag{25}$$

where $L_t^2 = L^2 + it/2m$. The probability density for finding the particle a distance r from the origin at time t is

$$|\psi(r,t)|^2 \propto \exp(-r^2/D_t^2) \tag{26}$$

where $D_t^2 = (L^4 + t^2/4m^2)/L^2$; hence the device's wave-function will have essentially constant half-width as long as $t \ll 2mL^2$. For a one-kilogram device whose initial half-width is 10^{-11}m, this means that the half-width will stay approximately constant on timescales of order 10^6 years. The probability of finding the particle in a given volume a distance r from the origin, compared with the probability of finding it in a similar-sized volume around the origin, is $|\psi(r,0)|^2/|\psi(0,0)|^2 = \exp(-r^2/L^2)$. If we take $r = 1\mu$m, then this probability ratio is approximately equal to $10^{-10^{10}}$. Only if we are working to 10^{10}-significant-figure accuracy do we need to allow for terms like this, and calculations are seldom done to such accuracy!

The truth, of course, is that the real state of any realistic model of NRQM is an entangled mess. It contains *no* exactly localised objects and *no* truly isolated subsystems, and there is *no* chance of working out its exact dynamics. Our actual strategy is to find a decomposition into subsystems such that they are *approximately* isolated, and to accept the errors that may result from this strategy.

We can quantify this by means of the Hilbert-space norm. Let $|\psi\rangle$ be the actual state of the system, and $\widehat{U}(t)$ the unitary operator generating its actual time evolution. Then to say that the state is approximately two isolated subsystems (during a period of length T) is to say that:

1. There is some state $|\sim\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle$ such that $|\langle\psi|\sim\psi\rangle| \simeq 1$;

2. There is some unitary operator $\widehat{U}_\sim(t) = \widehat{U}_1(t) \otimes \widehat{U}_2(t)$, such that, for $t < T$, $|\langle\psi| \widehat{U}^\dagger(t)\widehat{U}_\sim(t) |\sim\psi\rangle| \simeq 1$.[17]

(Note that we are not requiring $\widehat{U}_\sim(t) \simeq \widehat{U}(t)$ in any sense; our discussion is entirely state-dependent.)

If our standard description of a certain complex object (a DNA strand, a Geiger counter, a computer, whatever) describes it as an isolated object, and if isolated subsystems are only ever approximations to the 'true' quantum state, it might seem that we should say something like 'complex objects don't exist, they only approximately exist'. But this is to misunderstand the status of complex objects in physics. Such objects are actually identified by their structural properties (interacting with other DNA strands in a certain way, detecting radiation, calculating π, whatever) and such structural properties can tolerate small amounts of noise. (Put another way, an object which operates like a computer with 99.999999...% accuracy is still a computer!) From this viewpoint, to say that a certain complex object is present is to say that making that approximation is an extremely effective method of analysing the dynamics. (See Wallace (2001b) for a more extended discussion of this approach to complex objects).

Localisation enters this framework as a pragmatic criterion for isolation. If, in the state $|\sim\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle$, the states $|\psi_1\rangle$ and $|\psi_2\rangle$ have negligible amplitude to be in the same spatial location, and if the actual dynamics (as described by $\widehat{U}(t)$) are generated by a Hamiltonian in which the inter-particle forces drop off strongly with distance, then we can apply the isolated-subsystem approximation with considerable confidence as to its accuracy; conversely, if at some point this approximation predicts that $|\psi_1\rangle$ and $|\psi_2\rangle$ have

[17] There are two simplifications in this account. Firstly, in general a state of NRQM contains large numbers of identical particles, so the tensor product in (1) should be symmetrised and/or antisymmetrised to allow for these particles. Secondly, many subsystems (particularly those of macroscopic dimension) undergo decoherence, and hence cannot literally be said to be isolated in the sense used here. For systems whose dynamics are regular (i. e. non-chaotic) this is not particularly significant: decoherence will be significant if the system is not in some state of a certain preferred basis, but if it is in such a state then the isolated quantum dynamics will fairly accurately predict its evolution. For chaotic systems we cannot ignore decoherence in this manner (see Zurek (1998b) for details).

come to be localised, with high amplitude, in the vicinity of one another, then we should abandon the approximation from that point onwards.

This means that — not withstanding the existence in NRQM of a precisely definable position operator — we should treat the spatial location of a given object as approximate. It is not that my Geiger counter is localised in a number of exact locations with different amplitudes for each; rather, it is determinately[18] localised in some approximately-defined location. The exponentially small tails of the centre-of-mass wavefunction are totally irrelevant to any considerations of the counter's spatial location, for it is only in the approximation where these are neglected that we can treat the Geiger counter as isolated anyway.

Before applying these ideas to QFT, we should acknowledge a lacuna in the argument: we have assumed that the Hilbert-space norm is the appropriate measure of closeness of approximation, but what justifies this assumption? From a purely mathematical point of view, to be sure, it is an extremely obvious choice: it is a very natural metric on the space of states, and the unitarity of Hilbert-space dynamics means that the Hilbert-space distance between two points is time-invariant. But the situation is more complicated than this: Reutsche (1998) (writing in the context of the modal interpretation of quantum mechanics) points out that whilst a state like

$$\alpha \, |\text{Spin up}\rangle + \beta \, |\text{Spin down}\rangle \tag{27}$$

(where $\alpha \simeq 1$ and $\beta \simeq 0$) may legitimately be regarded as 'almost' a spin-up state, it is highly problematic to suppose that a macroscopic superposition like

$$\alpha \, |\text{Living cat}\rangle + \beta \, |\text{Dead cat}\rangle \tag{28}$$

is almost a live cat, whatever the values of α and β. She describes this sort of approximation (borrowing her terminology from Teller 1984) as "ontological distortion", as against the mere "numerical distortion" that occurs when, say, the position of a classical particle is predicted with slight inaccuracy: a macroscopic superposition is *not* (prima facie) even approximately the same as a system in a macroscopically definite state, but an ontologically new and problematic entity.

It is now clear that justifying the use of the Hilbert-space norm in this way takes us into the murky waters of the quantum measurement problem. We can see this another way by considering how easily an instrumentalist could dismiss our worries: from his perspective, the empirical content of the Hilbert-space norm is simply given by the probabilities of measurements to yield certain results, and states differing only very slightly in Hilbert space will give only slightly different probabilities for a given outcome on measurement. It

[18] No comment on the measurement problem is intended here; Everettians should read this as 'determinately localised relative to the branch which we are considering'.

is the illegitimacy of using "measurement" as a primitive concept in our approach which causes the difficulties.

In fact (staying in the non-relativistic domain) our problem is very closely related to some standard foundational questions in a number of interpretations:

- In the modal interpretation (van Fraassen 1991; Dieks and Vermaas 1998) the question is essentially equivalent to the Albert-Loewer (Albert and Loewer 1990, 1991; Albert 1992) problem of vagueness, discussed by Reutsche (1998) (see also references therein): the idea of a modal interpretation is to pick out a certain set of properties as definite, whilst it seems likely that the properties actually picked out are slightly different from the "right" ones — where 'slightly' means 'slightly in Hilbert-space norm'. The standard ('vague property') response to this is to replace precise properties (like spatial localisation) with a family of 'vague' properties, all very close in Hilbert-space norm to the 'precise' property under consideration.[19]

- In collapse interpretations such as those of Ghirardi, Rimini, and Weber (1986) and Pearle (1989) the question is equivalent to the 'problem of tails' (Albert and Loewer 1996): the collapse mechanism is generally chosen to cause the wave-function to be sharply peaked around a given spatial position but it is impossible to construct mechanisms which prevent the wave-function having exponential tails extending to infinity. The problem of tails can be addressed (Albert and Loewer 1996) by postulating some rule such as "a particle is localised in spatial region $\mathcal{R}$ if the squared modulus of the wave-function integrated over $\mathcal{R}$ exceeds $(1-p)$", where p is taken as small. There has been recent controversy (Lewis 1997; Ghirardi and Bassi 1999; Bassi and Ghirardi 1999, Clifton and Morton 1999, 2000) as to whether this rule should be understood as a metaphysical principle (which leads to difficulties both because of the arbitrariness of p and because of alleged failures of the arithmetic of macroscopic objects) or, as advocated by Clifton and Monton (1999), as a mere *facon de parler* useful in describing the wave-function (in which case there is presumably need for a defence of treating the tails as irrelevant to observational predictions).

- From the viewpoint of the Everett interpretation (at least in the form advocated by the author (Wallace 2001b; Wallace 2001d), which is essentially that developed by Zurek(most explicitly in Zurek 1998a) and Saunders (1995, 1996, 1998)) our problem

[19] In fact, the considerations of this section suggest a more serious problem for the modal interpretation: the very concept of macroscopic (non-relativistic) subsystems is threatened by wave-function spreading. Reutsche (1998) discusses a library book, localised approximately in her office but with exponential tails giving it a nonzero amplitude to be anywhere in space, and claims that such a book 'is in no way compatible with my ordinary notion of *located in my office*' ((Reutsche 1998, p. 233); italics hers) — but in fact, the interaction of those tails with the walls of the office will lead to the sort of rapid interaction that removes any prospect of distinguishing 'book' and 'wall' subsystems.

is that of justifying the neglect of parts of the wave-function of extremely low weight. This is essentially the probability problem of Everett interpretations: what justifies our strategy of disregarding low-weight branches when making decisions?

Although expanding on these rather brief comments would take us deep into interpretational questions and far beyond the scope of this paper, it seems likely that (at least for the interpretations above) a resolution of the quantum measurement problem would include a justification of the Hilbert-space norm as the appropriate measure of approximation. This seems particularly certain in the Everett interpretation, which is probably the interpretational program most naturally extendible to QFT and most in keeping with the spirit of this paper.

Given an interpretational justification of our use of the Hilbert-space norm, the concepts above carry through to QFT with only minor changes. The dynamical impossibility of constructing exactly isolated subsystems in NRQM becomes a kinematic impossibility in QFT if we consider the spatially localised subsystems $\mathcal{H}_{\Sigma_i}$, but for essentially the same reasons: interactions between those subsystems are so prevalent that the energy cost of constructing non-entangled states would take us out of the domain of validity of the QFT altogether. Also, in QFT it is even more apparent than in NRQM that objects can only be defined approximately unless they are isolated, for all that exist in QFT are the field observables, excitations of them, and patterns in those excitations, and if two patterns overlap and interact then there can be no exact criteria for individuating them. And in QFT there are no action-at-a-distance interactions, so spatial location is an even better criterion for isolation: two subsystems can be treated as isolated iff they are effectively localised in spatially separated regions.

Consequently, the approximate criteria developed above for NRQM apply just as well in QFT. In particular, QFT particles (unlike NRQM particles, which are part of the basic formalism) are themselves just certain patterns of excitations in the fields. As such, the appropriate concept of localisation for them is approximate, and so perfectly adequately treated by the methods of section 5.2. If we say, for instance, that a particle is localised in (spatial) region Σ_i, and that the region in the vicinity of Σ_i is empty, then what we mean is:

- that the field state restricted to Σ_i coincides almost exactly with the restriction to Σ_i of some state in the one-particle Hilbert space $\mathcal{H}_{1P}$;
- that in the region surrounding Σ_i the field state almost coincides with the vacuum.

These requirements are perfectly consistent with the actual field state being that of a particle (or of a particle together with some other excitations, created by operators strictly localised far from Σ_i); they are also consistent with the actual state itself being strictly (i. e., Knight-) localised in Σ_i. The point is, we neither know nor care: different ways of realising the requirements just mean different — but tiny — amounts of noise added to our highly-accurate description of the quantum state and of its dynamics. Ditto for particle *detectors*: there is no more need (and no more prospect) to regard them as exactly localised, provided that they are very well localised. If we model such a detector as made out of particles then that model will describe it as differing from the vacuum all the way to infinity, but it will also describe those differences as incredibly small beyond a certain very-well-defined, but not *precisely* defined, region. Then if the actual state of the QFT is almost the same as that of the model within that region, and is almost the same as the vacuum outside that region, then the criterion has been met for us to say that the region contains a detector.[20]

6 Conclusion

This completes our analysis of Lagrangian QFTs. We have argued that such QFTs can be made into perfectly well-defined quantum theories provided we take the high-energy cutoff absolutely seriously; that the multiple ways of doing this are not in conflict provided that we understand them as approximations to the structure of some deeper, as yet unknown theory; that the existence of inequivalent representations is not a problem; that a concept of localisation can be defined for such theories which is adequate to analyse at least some of the practical problems with which we are confronted; and that the inexactness inherent in that concept is neither unique to relativistic quantum mechanics, nor in any way problematic.

If there is one underlying theme to the approach to QFT advocated in this paper, it is this: the sort of information which we are interested in getting from physical theories is structural information. The reason that states with isolated subsystems approximate real quantum states well (both kinematically and dynamically) is because the two have virtually the same structure; the reason that we can be indifferent as to which concrete realisation of a scale-l AQFT to use is that structurally they are all virtually identical (at least until we get to scales of order l, at which point we no longer believe that any of them give us valid structural information).

The need to understand theories in structural terms is not restricted to QFT, but it is more starkly obvious there. It is not just that QFT is strictly speaking false, or that it is expected to be replaced someday by a deeper 'theory X'; that is true for virtually every physical theory we study. It is rather that — if Wilson's approach to renormalization is taken seriously — QFT only makes sense if we include in it some vestigial aspects of the

[20]We could, to be sure, define a detector as being present if the QFT state is almost the same as the model state within the localisation region, and *exactly* the same as the vacuum outside it — but what would be the justification for so doing? We don't have any reason to suppose that actual lab detectors are precisely localised in this way.

very theory which we expect to replace it.

From this viewpoint, we can see that Lagrangian QFT (as I have defended it) is not really in conflict with AQFT at all. Success in the AQFT program would leave us with a field theory exactly defined on all scales, and such a theory would be a perfectly valid choice for 'theory X': furthermore, even if we found such an exact QFT it would not prevent us from defining low-energy, 'effective' QFTs — which would not be well defined without a cutoff; nor, probably, would it obviate the need for these theories in describing certain low-energy limits of X.

If AQFT has any rival programs, in fact, they are string theory and other theory-of-everything candidates. Success in any of these programs would, of course, revolutionise physics, but that success would scarcely change the current status of Lagrangian QFT: as an inherently approximate, but nonetheless extraordinarily powerful tool to analyse the deep structure of the world.

Acknowledgements

I would like to thank Simon Saunders, whose own work on the foundations of QFT prompted my interest in the subject; and Jeremy Butterfield for many useful comments on an earlier draft of this paper. I have also benefitted greatly from conversations with Ian Aitchison, James Binney, Katherine Brading, Keith Burnett, Peter Morgan, and Andrew Steane.

References

Abraham, R. and J.E. Marsden. 1978. *Foundations of Mechanics.* 2nd Edition. Reading, Mass.: Benjamin/Cummings.

Albert, D. 1992. *Quantum mechanics and experience.* Cambridge, Massachussets: Harvard University Press.

Albert, D. and B. Loewer. 1990. "Wanted Dead or Alive: two attempts to solve Schrödinger's paradox." In *Proceedings of the 1990 Biennial Meeting of the Philosophy of Science Association*, edited by A. Fine, M. Forbes, and L. Wessels, Volume 1, 277–285. East Lansing, Michigan: Philosophy of Science Association.

———. 1991. "The measurement problem: some 'solutions'." *Synthese* 86:87–98.

———. 1996. "Tails of Schrödinger's cat." In *Perspectives on Quantum Reality*, edited by R. Clifton, 81–92. Dordrecht: Kluwer Academic Publishers.

Barbour, J. B. 1994a. "The timelessness of quantum gravity: I. The evidence from the classical theory." *Classical and Quantum Gravity* 11:2853–2873.

——— 1994b. "The timelessness of quantum gravity: II. The appearence of dynamics in static configurations." *Classical and Quantum Gravity* 11:2875–2897.

——— 1999. *The End of Time*. London: Weidenfeld and Nicholson.

Bassi, A. and G. Ghirardi. 1999. "More about dynamical reduction and the enumeration principle." *British Journal for the Philosophy of Science* 50:719–734. Available online at http://xxx.arXiv.org/abs/quant-ph/9907050.

Bennett, C. H., G. Brassard, R. Jozsa, A. Peres, and W. K. Wootters. 1993. "Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels." *Physical Review Letters* 70:1895–1898.

Binney, J. J., N. J. Dowrick, A. J. Fisher, and M. E. J. Newman. 1992. *The Theory of Critical Phenomena : an introduction to the renormalisation group*. Oxford: Oxford University Press.

Buchholz, D. 2000. "Algebraic Quantum Field Theory: a Status Report." Plenary talk given at the XIIIth International Conference on Mathematical Physics, London. Available online at http://xxx.arxiv.org/math-ph/0011044.

Cao, Tian Yu. 1997. *Conceptual Developments of 20th Century Field Theories*. Cambridge: Cambridge University Press.

Carlip, S. 1998. *Quantum Gravity in 2+1 Dimensions*. Cambridge: Cambridge University Press.

Castellani, E. 2000. "Reduction, emergence, and effective field theories." Forthcoming in *Studies in the History and Philosophy of Modern Physics*. Available online from http://philsci-archive.pitt.edu/.

Clifton, R. and H. Halvorson. 2001. "Entanglement and Open Systems in Algebraic Quantum Field Theory." *Studies in the History and Philosophy of Modern Physics* 32:1–31. Available online at http://xxx.arXiv.org/abs/quant-ph/0001107.

Clifton, R. and B. Monton. 1999. "Losing your marbles in wavefunction collapse theories." *British Journal for the Philosophy of Science* 50:697–717. Available online at http://xxx.arXiv.org/abs/quant-ph/9905065.

———. 2000. "Counting marbles with 'accessible' mass density: a reply to Bassi and Ghirardi." *British Journal for the Philosophy of Science* 51:155–164. Available online at http://xxx.arXiv.org/abs/quant-ph/9909071.

Deutsch, D. 1984. "Towards a Quantum Theory without 'Quantization'." In *Quantum Theory of Gravity: essays in honour of Bryce S. DeWitt*, edited by S.M. Christensen, 421–430. Bristol: Adam Hilger Ltd.

———. 1985. "Quantum Theory as a Universal Physical Theory." *International Journal of Theoretical Physics* 24 (1): 1–41.

Deutsch, D. and P. Hayden. 2000. “Information Flow in Entangled Quantum Systems.” *Proceedings of the Royal Society of London* A456:1759–1774. Available online at http://www.arXiv.org/abs/quant-ph/9906007.

Dieks, Dennis and P. E. Vermaas, eds. 1998. *The Modal Interpretation of Quantum Mechanics.* Dordrecht: Kluwer Academic Publishers.

Enderton, H. B. 1977. *Elements of Set Theory.* San Diego, California: Academic Press.

Fredenhagen, K. 1985. “A remark on the cluster theorem.” *Communications in Mathematical Physics* 97:461–463.

Ghirardi, G. and A. Bassi. 1999. “Do dynamical reduction models imply that arithmetic does not apply to ordinary macroscopic objects?” *British Journal for the Philosophy of Science* 50:705–720. Available online at http://xxx.arXiv.org/abs/quant-ph/9810041.

Ghirardi, G., A. Rimini, and T. Weber. 1986. “Unified dynamics for micro and macro systems.” *Physical Review D* 34:470–491.

Haag, R. 1996. *Local Quantum Theory: Fields, Particles, Algebras.* Berlin: Springer-Verlag.

Halvorson, H. and R. Clifton. 2001. “No place for particles in relativistic quantum theories?” Available online at http://xxx.arXiv.org/abs/quant-ph/0103041.

Hartmann, S. 2000. “Effective field theories, reductionism and scientific explanation.” Forthcoming in *Studies in the History and Philosophy of Modern Physics.* Available online from http://philsci-archive.pitt.edu/.

Hegerfeldt, G. A. 1998a. “Causality, Particle Localization and Positivity of the Energy.” In *Irreversibility and Causality*, edited by A. Böhm, 238–245. New York: Springer.

——— 1998b. “Instantaneous spreading and Einstein causality.” *Annalen der Physik* 7:716–725.

Knight, J. M. 1961. “Strict localisation in quantum field theory.” *Journal of Mathematical Physics* 2:459.

Ladyman, J. 1998. “What is Structural Realism.” *Studies in the History and Philosophy of Science* 29:409–424.

Lewis, P. J. 1997. “Quantum Mechanics, Orthogonality, and Counting.” *British Journal for the Philosophy of Science* 48:313–328.

Marsden, J. E. and T. S. Ratiu. 1994. *Introduction to Mechanics and Symmetry.* New York: Springer-Verlag.

Morrison, H. L. 1962. *The Quantum Theory of Many-Particle Systems.* New York: Gordon and Breach Science Publishers.

Newton, T. D. and E. P. Wigner. 1949. "Localized states for elementary particle physics." *Reviews of Modern Physics* 21:400.

Pearle, P. 1989. "Combining stochastic dynamical state-vector reduction with spontaneous localization." *Physical Review A* 39 (5): 2277–2289.

Peskin, M. E. and D. V. Schroeder. 1995. *An introduction to Quantum Field Theory.* Reading, Massachusetts: Addison-Wesley.

Psillos, S. 1995. "Is Structural Realism the Best of Both Worlds?" *Dialectica* 49:15–46.

Redhead, M. 1995. "More Ado about Nothing." *Foundations of Physics* 25:123.

Reeh, H. and S. Schlieder. 1961. "Bemerkungen zur Unitäräquivalenz von Lorentzinvarianten Feldern." *Nuovo Cimento* 22:1051.

Reutsche, L. 1998. "How Close is 'Close Enough'?" In Dieks and Vermaas 1998, 223–240.

Saunders, S. 1995. "Time, Decoherence and Quantum Mechanics." *Synthese* 102:235–266.

——— 1996. "Time, Quantum Mechanics and Tense." *Synthese* 107:19–53.

——— 1998. "Time, Quantum Mechanics, and Probability." *Synthese* 114:373–404.

Segal, I. 1964. "Quantum Fields and Analysis in the Solution Manifolds of Differential Equations." Edited by W. Martin and I. Segal, *Analysis in Function Space.* Cambridge, Massachusets: MIT Press.

Strocchi, F. and A. S. Wightman. 1976. *Journal of Mathematical Physics* 17:1930.

Teller, P. 1984. "The Projection Postulate: a new perspective." *Philosophy of Science* 51:369–395.

van Fraassen, B. C. 1991. *Quantum Mechanics.* Oxford: Oxford University Press.

Wald, R. M. 1984. *General Relativity.* Chicago: University of Chicago Press.

——— 1994. *Quantum Mechanics on Curved Space-Time and Black Hole Thermodynamics.* Chicago: University of Chicago Press.

Wallace, David. 2001a. "Emergence of particles from bosonic quantum field theory." Available online from http://phil-sci.pitt.edu.

——— 2001b. "Everett and Structure." To appear in *Studies in the History and Philosophy of Modern Physics*; available online at http://xxx.arXiv.org/abs/quant-ph/0107144 or from http://phil-sci.pitt.edu.

——— 2001c. "Implications of Quantum Theory in the Foundations of Statistical Mechanics." Available online from http://philsci-archive.pitt.edu.

——— 2001d. “Worlds in the Everett Interpretation.” Forthcoming in *Studies in the history and philosophy of modern physics*; available online at http://www.arxiv.org/abs/quant-ph/0103092 or from http://phil-sci.pitt.edu.

Wilson, K. G. and J. Kogut. 1974. “The Renormalization Group and the ϵ Expansion.” *Physics Reports*, vol. 12C.

Woodhouse, N. M. J. 1991. *Geometric Quantization.* Oxford: Oxford Science Publications.

Worrall, J. 1989. “Structural Realism: the Best of Both Worlds.” *Dialectica* 43:99–124. Reprinted in *The Philosophy of Science*, D. Papineau (Ed.), Oxford University Press, 1996, pp. 139–165.

Zurek, W.H. 1998a. “Decoherence, Einselection and the Existential Interpretation (The Rough Guide).” *Philosophical Transactions of the Royal Society of London* A356:1793–1820.

——— 1998b. “Decoherence, Chaos, Quantum-Classical correspondence, and the algorithmic arrow of time.” *Phys. Scripta* T76:186–198. Available online at http://xxx.arXiv.org/abs/quant-ph/9802054.