

Spontaneous Symmetry Breaking in Finite Quantum Systems: a decoherent-histories approach

David Wallace*

August 28, 2018

Abstract

Spontaneous symmetry breaking (SSB) in quantum systems, such as ferromagnets, is normally described as (or as arising from) degeneracy of the ground state; however, it is well established that this degeneracy only occurs in spatially infinite systems, and even better established that ferromagnets are not spatially infinite. I review this well-known paradox, and consider a popular solution where the symmetry is explicitly broken by some external field which goes to zero in the infinite-volume limit; although this is formally satisfactory, I argue that it must be rejected as a physical explanation of SSB since it fails to reproduce some important features of the phenomenology. Motivated by considerations from the analogous classical system, I argue that SSB in finite systems should be understood in terms of the approximate decoupling of the system's state space into dynamically-isolated sectors, related by a symmetry transformation; I use the formalism of decoherent histories to make this more precise and to quantify the effect, showing that it is more than sufficient to explain SSB in realistic systems and that it goes over in a smooth and natural way to the infinite limit.

1 A story about symmetry breaking

There is a standard account of spontaneous symmetry breaking (SSB) in quantum theory that is physically intuitive, explanatorily powerful, and predictively effective, and it goes like this. To begin with, consider a single nonrelativistic particle of mass m , moving in one dimension in a smooth potential $V(x)$. In full generality, very little can be said about this particle's dynamics. But if we now suppose that the potential has a single minimum at $x = x_0$, we can expand the potential around that minimum:

$$V(x_0 + \eta) = V(x_0) + \frac{1}{2}V''(x_0)\eta^2 + O(\eta^3). \quad (1)$$

*School of Philosophy, University of Southern California; email dmwallac@usc.edu

And if the potential is sufficiently slowly changing, then at least for the lowest-lying energy eigenstates, we can neglect the terms in $(x - x_0)^n$ for $n > 2$ and approximate the system as a harmonic oscillator, with energy levels

$$E_n = V(x_0) + \hbar(n + 1/2)\omega. \quad (2)$$

where $\omega^2 = V''(x_0)/m$. The well-known methods of perturbation theory now allow us to make successive corrections to these energy levels due to the higher-order terms in the expansion of $V(x)$ around $x = x_0$, and also to determine when the harmonic oscillation starts to break down entirely.

Now suppose the potential is symmetric under $x \rightarrow -x$. If it has a single minimum at $x = 0$, the analysis above applies unchanged, with the extra detail that the perturbative corrections are also symmetric under this transformation, and so the structure of the spectrum of energy eigenstates shows that symmetry: in particular, since the spectrum is nondegenerate, each eigenstate is individually symmetric or antisymmetric with respect to the transformation. (And so in particular, the position-space wavefunction ψ of an eigenstate satisfies $\psi(-x) = \pm\psi(x)$.)

If instead there is a minimum at $x_0 > 0$, there will be another minimum at $x = -x_0$. For definiteness, we can consider the potential

$$V(x) = \frac{\lambda}{4!}x^4 - \frac{\mu^2}{2}x^2 + \frac{3\mu^4}{2\lambda} \quad (3)$$

with minima at $x = \pm x_0 \equiv \pm(6\mu^2/\lambda)^{1/2}$ (and where the constant is just for convenience, to enforce $V(x) = 0$ at the minima). Expanding around, say, the minimum at $x = +x_0$, we get

$$V(x_0 + \eta) = \frac{1}{2}(2\mu)^{1/2}\eta^2 + O(\eta^3) \quad (4)$$

which — if the minima are sufficiently widely separated, and the function's slope is sufficiently gentle — we can again treat as a perturbed harmonic oscillator. So at least for low-lying energy eigenstates, the system can (apparently) be treated as degenerate, with two ground states localised around $x = \pm x_0$ and two towers of oscillations, one built on each ground state. There will be perturbative corrections due to the higher-order terms in the expansion, and non-perturbative corrections due to quantum tunnelling between the minima, but (again, apparently) these have only minor effects on the physics provided the system is not too excited. However, the expansion of the potential around a given minimum will not be symmetric under $\eta \rightarrow -\eta$. So long as we confine our attention to one set of low-lying states, the symmetry will not be manifest. It will instead be detectable only in certain relations between the higher-order coefficients in the expansion.

In particular, at sufficiently low temperatures the statistical mechanics of the system will display a certain asymmetry despite the symmetry in $V(x)$: the statistical-mechanical state of the system will be *either* a canonical-ensemble state built out of the excitations localised around the left-hand potential well, *or*

the equivalent state built out of the excitations localised around the right-hand well. As the temperature is increased, this will break down, and eventually (and, for this small system, in a rather fuzzy and imprecise way) the energy eigenstates making a significant contribution to the canonical ensemble will be spread across both wells.

Qualitatively similar phenomena occur when we move from a single quantum particle to a quantum field theory¹ — say, to a scalar theory defined (either through canonical quantization, or directly via the path integral) by Lagrangian

$$\mathcal{L} = \int dx^3 \left(\frac{1}{2} \partial_\mu \phi(x) \partial^\mu \phi(x) - V(\phi(x)) \right) \quad (5)$$

where V is again an arbitrary smooth potential and where the field theory is defined in some *finite* region, say of volume L^3 (and where we use units with $c = 1$). If V has a single minimum at $\phi = \phi_0$, we can expand around that minimum —

$$V(\phi_0 + \eta) = V(\phi_0) + \frac{1}{2} V''(\phi_0) \eta^2 + O(\eta^3) \quad (6)$$

— and treat the field theory, to first approximation, as a free field theory with Lagrangian

$$\mathcal{L}_0 = \int dx^3 \left(\frac{1}{2} \dot{\eta}(x)^2 - \frac{1}{2} (\nabla \eta(x))^2 - m^2 \eta(x)^2 \right). \quad (7)$$

The ground state $|\Omega\rangle$ of that theory — the *vacuum*, in particle-physics terminology — can be treated — again, to first approximation — as the ground state of a collection of uncoupled oscillators corresponding to the modes of the free field; its excitations, one family per mode, are characterised by a *dispersion function* $\omega(k)$, so that $\hbar\omega(k)$ is the energy of an excitation of momentum $\hbar k$. In our particular example, $\hbar\omega(k) = (m^2 + \hbar^2 k^2)^{1/2}$; in more general theories (in particular, in non-Lorentz-covariant field theories like those encountered in solid-state physics) the dispersion function can have a more general form. In any case, the excitations can be interpreted (at least approximately) as *particles*, and the non-quadratic corrections to $V(\phi_0 + \eta)$ can be treated perturbatively as interactions between the particles, giving rise to scattering. The expectation value of the field operator on the vacuum is

$$\langle \Omega | \widehat{\phi}(x) | \Omega \rangle = \phi_0. \quad (8)$$

(This whole framework requires some kind of short-distance cutoff to be mathematically well-behaved, and the exact issue of the relation between field and particle has its own subtleties, but these issues are orthogonal to those I want to consider here; see Wallace (2011, 2017) and references therein for further discussion.)

¹For details on quantum field theory as it is used in high-energy physics, see (e.g.) Peskin and Schroeder (1995), Weinberg (1995a, 1995b), Duncan (2012); for quantum field theory as it is used in solid-state physics, see Altland and Simons (2010); for a conceptual introduction aimed at non-specialists, see Wallace (2017).

If the potential function has a minimum at $\phi = 0$ and is symmetric under $\phi \rightarrow -\phi$, that symmetry will be manifest in the particle structure of the theory and in particular in the multiparticle interactions. If the potential is symmetric but has minima away from $\phi = 0$ — for definiteness again, if it has the quartic form (3) — we expect *two* vacuum states $|\Omega+\rangle, |\Omega-\rangle$, with field expectation values

$$\langle \Omega\pm | \hat{\phi}(x) | \Omega\pm \rangle = \pm\phi_0. \quad (9)$$

The excitations around each vacuum can be described in particle language again, but the symmetry is now hidden, and cannot be seen explicitly in multiparticle interactions.

If we consider the statistical mechanics of this theory, at low energies we expect again two different thermal equilibria, each consisting in the free-particle approximation of a gas of excitations around a given vacuum. As the temperature increases the relevant excitations will begin to spread out of the separate potential wells, so that eventually the symmetry will be restored. Depending on the system it is at least plausible that this will constitute a *phase transition*, characterised by changes in the functions of state or their derivatives which become sharper and sharper as L gets larger, tending smoothly in the limit as $L \rightarrow \infty$ to discontinuous change.

This somewhat heuristic and perturbative description of SSB in field theory can be put on a firmer footing using the *effective action* formalism (Coleman and Weinberg 1973; Weinberg 1973; Dolan and Jackiw 1974). Its details will not be salient for this paper (for a pedagogical introduction in the zero-temperature regime, see Peskin and Schroeder (1995, ch.11)), but it allows us to define a temperature-dependent functional — the “effective action” —

$$\Gamma[\phi; T] \quad (10)$$

of a classical field ϕ with the properties that:

1. The effective action has the same symmetries as the underlying field theory.
2. If ϕ_c solves

$$\left. \frac{\delta\Gamma[\phi; T]}{\delta\phi} \right|_{\phi=\phi_c} = 0 \quad (11)$$

then $\phi_c(x)$ is the expectation value of $\hat{\phi}(x)$ in the thermal-equilibrium state at temperature T . In particular, at $T = 0$, $\phi_c(x)$ is the vacuum expectation value of $\hat{\phi}(x)$.

3. The eigenfunctions of the operator

$$\left. \frac{\delta^2\Gamma[\phi; T]}{\delta\phi(x)\delta\phi(y)} \right|_{\phi=\phi_c} \quad (12)$$

(understood as an operator on real functions on \mathfrak{R}^4) determine the elementary excitations of the field theory at temperature T , and in particular the mass of the particles.

So the thermal-equilibrium state of the field theory breaks symmetry at temperature T iff $\Gamma[\phi; T]$ is extremised for nonzero ϕ (and in this circumstance there are multiple extrema, and so multiple thermal equilibria). In particular the ground state is degenerate if $\Gamma[\phi; 0]$ has multiple extrema. If symmetry is not spontaneously broken then the symmetries of the effective action are also symmetries of (12) and hence of the particle spectrum. If it is spontaneously broken, the particle spectrum does not display the symmetry. And the signature of a phase transition is a theory where the extrema of the effective action are nonzero below some critical temperature T_c and zero above that temperature.

Finally, if the field theory has a (global) symmetry that is continuous, spontaneous breaking of that symmetry leads to *Goldstone bosons*: elementary excitations whose dispersion relations satisfy $\omega(0) = 0$. (In relativistic theories this corresponds to massless particles.) One example can be given just by taking our scalar field to be complex-valued, with Lagrangian

$$\mathcal{L} = \int dx^3 \left(\frac{1}{2} \partial_\mu \phi^*(x) \partial^\mu \phi(x) - V(|\phi(x)|) \right) \quad (13)$$

so that it is invariant under $\phi \rightarrow e^{i\theta} \phi$. For another illustrative example² consider nonrelativistic physics with finite particle density. At low temperatures, the equilibrium state will usually be some kind of crystal, and hence spontaneously breaks translation symmetry; the associated Goldstone bosons are the *phonons*, quanta of vibration, which have a dispersion relation satisfying $\lim_{k \rightarrow 0} \omega(k) = 0$. The statement that Goldstone bosons are a universal consequence of spontaneous breaking of a (global) continuous symmetry is *Goldstone's theorem*; it can be proved via the effective action by consideration of (12) and through several other ways³, and has as its central premise the existence of a vacuum state in which symmetry is broken.

To repeat: this account is physically intuitive, explanatorily powerful, and predictively effective. So it is a little disturbing to note that it cannot, strictly speaking, be correct.

2 Non-degeneracy of the ground state

Consider again the example of a single particle in a double-well potential, and for the moment imagine modifying that potential to introduce an infinitely-high potential barrier at $x = 0$. (Formally, we can replace $V(x)$ with

$$V'(x) = V(x) + K\delta(x) \quad (14)$$

where δ is the Dirac delta function, and take the limit as $K \rightarrow \infty$.) In this case, the dynamics entirely decouples into left-side and right-side dynamics. On the

²See Anderson (1984, ch.2) and references therein; or, for a more elementary introduction, Ashcroft and Mermin (1976, ch.23).

³See, e.g., Weinberg (1995b, section 19.2), Witten (2000, p.1136-1140), and references therein.

left, we will have an infinite tower of energy eigenstates $\psi_{L,n}$ —

$$\hat{H}\psi_{L,n}(x) = E_n\psi_{L,n}(x) \quad (15)$$

— with $\psi_{L,n}(x) = 0$ for $x > 0$. To first approximation these eigenstates will be harmonic-oscillator eigenstates with energy $E_n = \hbar(n+1/2)(V''(-x_0)/m)$. The non-quadratic part of the potential will perturb them away from this simple form and make $O(\hbar^2)$ corrections to the energy; in principle these perturbations will include the effect of the infinite barrier but in practice the details of this barrier will make very little difference for small n since $\psi_{L,n}(x)$ will be constrained to be very small for $x \simeq 0$ in any case.

On the right, there will be a similar family of eigenstates $\psi_{R,n}$. Given the symmetry, we have $\psi_{R,n}(x) = \psi_{L,n}(-x)$, and in particular $\psi_{R,n}(x) = 0$ for $x < 0$; the symmetry also entails that the eigenvalue of $\psi_{R,n}$ is also E_n . So we have exact degeneracy not only of the ground state but of all excitations: energy eigenstates appear in pairs.

Now take out the infinite potential barrier, so that the $\psi_{L,n}$ states are confined on the left side only by the finite potential barrier between the two wells. It is an elementary result of quantum mechanics that no finite potential can *exactly* confine a wavefunction. It may take a long time, but eventually any such wavefunction will tunnel through to the right-hand side. Which is to say that $\psi_{L,n}$ is not an eigenstate of energy — hence, not a stationary state — once the barrier is removed.

To explore this further: if we take symmetric or antisymmetric combinations of the $\psi_{L,n}$ and $\psi_{R,n}$, we expect those states to be very close to invariant under tunnelling, even for arbitrarily long times. So those symmetric and antisymmetric combinations are — to a very good approximation — the eigenstates of energy, but there will be an energy gap between them, generated by tunnelling effects and completely lifting the ground-state degeneracy. Or, to put it another way, we expect that (for sufficiently low-lying energies) the exact eigenstates of energy come in symmetric and antisymmetric pairs $\psi_{\pm,n}$ satisfying $\psi_{\pm,n}(x) = \pm\psi_{\pm,n}(-x)$, such that the combinations $\psi_{+,n} \pm \psi_{-,n}$ are very close to our original states $\psi_{L,n}$, $\psi_{R,n}$.

This heuristic argument can be established by direct calculation of the energy gap, using instanton methods. See Coleman (1985, ch.7) for the details, but in outline: the low-lying energy states, and in particularly the ground state, do come in pairs $\psi_{\pm,n}(x)$, with energies

$$E_{\pm,n} = (n+1/2)\hbar\omega + O(\hbar^2) \pm \frac{\kappa\hbar\omega}{2}(S/\hbar)^{1/2}e^{-S/\hbar}. \quad (16)$$

Here, the first term is the ordinary harmonic-oscillator energy. The $O(\hbar^2)$ term represents perturbative corrections to the harmonic oscillator which are the same for $E_{+,n}$ and $E_{-,n}$. And the exponential term gives the energy-level splitting. S itself is given by

$$S = \int_{-x_0}^{+x_0} dx \sqrt{2mV(x)} \quad (17)$$

and can be interpreted as the action of the *instanton* connecting the two classical minima $\pm x_0$. (κ is a dimensionless constant of order unity, whose exact form will not be relevant to us.) The splitting term is in general extremely small — and, in particular, it tends to zero as $\hbar \rightarrow 0$ faster than any power of \hbar , and so faster than any perturbative correction — but it is not zero, and that is enough to break the degeneracy.

Essentially the same effects occur in quantum field theory with a double-well potential. Again, given an infinite potential barrier between the two wells, we expect exact degeneracy, and two families of excitations — interpretable as particles — built one upon each ground state. But again, if the barrier is only finitely high then tunnelling can occur between the two potential wells and the degeneracy is lifted. The energy gap between (low-lying) pairs of states is given (via a straightforward generalisation of the method of instantons) by

$$\Delta E = \frac{\kappa m}{2} (S/\hbar)^{1/2} e^{-S/\hbar} \quad (18)$$

where now

$$S = L^3 \int_{-\phi_0}^{+\phi_0} d\phi \sqrt{2V(\phi)}. \quad (19)$$

The case of (apparently) spontaneously broken continuous symmetry is qualitatively similar but differs in the details. No energy barrier confines the system to any particular value of the average field strength, so any localised state will spread out in the same way that a wave-packet state of a free particle spreads out. In the case of the complex scalar field in particular, if we write $\phi = \rho \exp(i\theta)$ then the Lagrangian becomes

$$\mathcal{L} = \int dx^3 \left(\frac{1}{2} \rho(x)^2 \partial_\mu \theta(x) \partial^\mu \theta(x) + \frac{1}{2} \partial_\mu \rho(x) \partial^\mu \rho(x) - V(\rho(x)) \right). \quad (20)$$

In the approximation where interactions can be neglected, we can decouple the θ and ρ modes and get

$$\mathcal{L}_\theta = \frac{1}{2} \int dx^3 \rho_0^2 \partial_\mu \theta(x) \partial^\mu \theta(x) \quad (21)$$

where ρ_0 is the value of ρ that minimises $V(\rho)$. (This is one form of the *non-linear sigma model*; see, e.g., Weinberg (1995b, ch.19), Donoghue, Golowich, and Holstein (2014, ch.IV-1) and references therein.) If we define

$$\bar{\theta} = \frac{1}{L^3} \int dx^3 \theta(x) \quad (22)$$

and write $\theta(x) = \bar{\theta} + \delta\theta(x)$, then the Lagrangian decouples further:

$$\mathcal{L}_\theta = \frac{\rho_0^2 L^3}{2} \dot{\bar{\theta}}^2 + \frac{1}{2} \int dx^3 \rho_0^2 \partial_\mu \delta\theta(x) \partial^\mu \delta\theta(x). \quad (23)$$

The average value of the phase just has the Lagrangian for a free particle with mass $(\rho_0^2 L^3)^{-1}$, moving on the unit circle; clearly, when quantized this system

has a nondegenerate ground state with wavefunction $\Psi[\bar{\theta}] = (2\pi)^{-1/2}$. Similarly, given a finite-size crystal, the wavefunction of the center of mass will spread out freely until it is uniformly spread over the region accessible to the crystal, so again the ground state, which is time-independent, must be nondegenerate. (For a similar analysis of the harmonic crystal, see van Wezel and van den Brink (2007).)

So: in all the systems we considered in section 1, the lowest-energy state is nondegenerate. (And, although the arguments I have given involve various approximation schemes, this can be proven rigorously for one-particle systems and at least a wide class of field theories; see Witten (2000, pp.1125-1130) for details.) And this seems to make Section 1's analysis, which relied on ground state degeneracy, completely untenable.

3 The return of degeneracy in infinite systems

On a purely formal level, degeneracy can in a sense be recovered by considering *infinite* quantum systems, using the tools of algebraic quantum field theory, or AQFT.⁴ In AQFT, a quantum field theory is characterised in the first instance by the algebraic structure of the field observables, without stipulating in advance a representation of those field observables as Hilbert-space operators. Particular models of the quantum field theory then correspond to particular representations of those observables as operators on some Hilbert space. This algebraic approach is in principle available for ordinary quantum-mechanical systems, but in practice adds little since (under weak assumptions; cf Ruetsche (2011, ch.3)) any two such representations are unitarily equivalent, i. e. are related by some unitary transformation. In infinite-volume quantum field theories⁵ this is not the case: there are many (indeed, uncountably many) unitarily inequivalent representations. In particular, we can understand SSB in this framework as the existence of inequivalent representations, differing by the spatially-averaged value of the field whose expectation value breaks the symmetry. In each such representation the ground state remains nondegenerate, but the symmetry transformation is represented not by a unitary transformation *of* each representation, but by a non-unitary map *between* representations.

But of course, real condensed-matter systems are manifestly *not* infinite; *perhaps* particle-physics systems should be thought of as infinite, but they are at any rate not fields in infinite *Minkowski* spacetime, as they are represented in AQFT. In both cases, the spatially-infinite description seems something that we should regard as a useful idealisation or approximation of some finite system (or some finite patch of a larger cosmological system) and not something to be

⁴For reviews of AQFT, see, e. g. , Haag (1996), Ruetsche (2011, ch.1-5), or Halvorson (2007); for conceptual discussions of SSB from an AQFT viewpoint, see Ruetsche (2006, 2011 ch.12-14) and Baker and Halvorson (2013).)

⁵And also in finite-volume quantum field theories without the short-distance cutoff I introduced in section 1, though it is doubtful whether this has any physical significance; see Wallace (2011, 2017) for further discussion of this point.

taken literally.⁶ From this perspective (urged upon us by Earman (2004) and Butterfield (2011)), physical phenomena that occur in the infinite limit ought to be taken seriously only when they are in some sense approximated by physical phenomena that occur in finite systems, with the approximation getting better as the size of the finite system increases. We see this in phase transitions: their strict mathematical characterisation is in terms of a discontinuity in their state functions or the derivatives of those functions; this discontinuity will not occur in any finite system; but the function will change increasingly rapidly and will tend towards a discontinuous function as the system volume tends towards infinity. (See Butterfield (2011) for further discussion of this example.)

It is not at all clear how to understand SSB in this way. The ground state of our systems remains resolutely non-degenerate, however big the system is, and so is invariant under the symmetry operation. So the expectation value of the supposedly symmetry-breaking field, with respect to the ground state, must remain exactly zero no matter how big the system is, and does not in any obvious sense tend towards the nonzero value that is obtained in an infinite system and that is calculated perturbatively, and empirically successfully, using effective-action methods. Nor can the effective action even really be defined for infinite systems: it is proportional to the spatial volume of the system and so is infinite for infinite systems.

To summarise: a naive consideration of symmetry breaking in finite systems leads to the prediction of a multiplicity of ground states transforming nontrivially under the symmetry, a degenerate collection of excitations built on that ground state, and associated phenomenology (such as symmetry restoration at high temperatures, and Goldstone bosons). Those predictions are confirmed empirically, to great accuracy, and yet the naive considerations they rest on are *false*, and the ground state of finite systems shows no degeneracy. Degeneracy does occur in genuinely infinite systems, and those genuinely infinite systems reproduce at least some of the phenomenology of SSB. But the temptation to regard those infinite systems as the correct description of the physics (a position discussed sympathetically by Ruetsche (2003, 2011) and Batterman (2001, 2005)) is at least *prima facie* in conflict with the obvious fact that solid-state systems are of finite size, and the (somewhat less obvious) fact that particle-physics systems are only accurately described as fields on Minkowski spacetime when confined to finite regions; it is also in conflict with empirically-very-powerful calculational tools like the effective action, which apparently require finite volume to be well defined.

In the remainder of the paper I will explore how SSB can, after all, be understood in finite systems. But in this exploration it will be important to remember that the infinite limit cannot be simply disregarded. Infinite-system

⁶Feynman: “[I]t is not often that experiments are done under the stars. Rather they are done in a room. Although it is physically reasonable that the walls have no effect, it is true that the original problem is set up as an idealisation. It is no more satisfactory idealization to remove the walls to infinity than to replace them by perfect mirrors far away. The mathematical rigor is wasted in the first idealization, since the walls are not at infinity.” (Feynman and Hibbs 1965, pp.93-4).

models do succeed in describing SSB, and it is implausible that this is a coincidence: our task is not just to understand the process in finite systems but also to understand in what sense those systems' behaviour tends towards the symmetry-breaking behaviour of the infinite systems in the limit of large size.

4 A partial solution: symmetry breaking from the environment

Suppose we return to the double-well potential (1) and add a linear symmetry-breaking term:

$$V(x) \rightarrow V_J(x) = V - Jx \quad (24)$$

(with, for definiteness, $J > 0$). The effect will be to explicitly break the symmetry: to a first approximation, the classical ground state at $x = +x_0$ will decrease in energy by Jx_0 and the ground state at $x = -x_0$ will increase by a similar amount. (The linear term will also slightly displace the minimum of the potential.) Provided that Jx_0 is large compared to the tunnelling-induced energy gap $\kappa\hbar\omega(S/\hbar)^{1/2}e^{-S/\hbar}$ (as defined by (17)) between the pairs of energy eigenstates, this term will drastically change the shape of the energy eigenstates: now the ground state will be very close to a harmonic-oscillator ground state at $x \simeq +x_0$, and there will be an excited state very close to a harmonic-oscillator ground state at $x \simeq -x_0$. Indeed, if $Jx_0 \ll \hbar\omega$, the excitation energy required to add a quantum of energy to either potential well, then this will be the *first* excited state. And, for at least the lower-lying excited states, the same will occur: the spectrum will continue to consist of closely-spaced pairs of eigenvalues, but the eigenstates of the eigenvalue pair around $(n + 1/2)\hbar\omega$ will be very close to the n th excited states of harmonic oscillator potentials localised at $\pm x_0$.

So the introduction of even this very small asymmetry restores the picture of the double-well potential as describing a system which can be localised in either one of two harmonic-oscillator potentials. The only difference between this picture and the picture of section 1 is that there is not exact degeneracy between the two oscillators: the family of excitations of the left-hand oscillator is higher in energy than that of the right-hand oscillator by $\simeq 2Jx_0$. (As usual, for sufficiently excited states this account will begin to break down as the states cease to be localised even approximately in one well or the other.)

Essentially the same story plays out for the continuum version of the double well potential, though the effect is more dramatic: we modify the Lagrangian (5) to

$$\mathcal{L}_J = \int dx^3 \left(\frac{1}{2} \partial_\mu \phi(x) \partial^\mu \phi(x) - [V(\phi(x)) - J\phi(x)] \right) \quad (25)$$

which will localise the ground state (in field-space, not physical space) around $\phi = +\phi_0$ provided that $J\phi_0$ is large compared to the energy gap given by (18). Since this energy gap decreases exponentially with system volume, as that volume gets larger the perturbing field can become smaller and smaller while still localising the ground state.

In for the case of continuous symmetry breaking, the appropriate modification is to add to the potential a term $-J\text{Re}(\phi)$ (that is: proportional to the real part of ϕ). The part of the Lagrangian controlling the average phase is then modified to

$$\mathcal{L}_{\bar{\theta}} = \frac{\rho_0^2 L^3}{2} \dot{\bar{\theta}}^2 + J\rho_0 L^3 \cos \bar{\theta} \quad (26)$$

and for small $\bar{\theta}$ this is just a harmonic-oscillator potential:

$$\mathcal{L}_{\bar{\theta}} \simeq \frac{\rho_0^2 L^3}{2} \dot{\bar{\theta}}^2 - \frac{J\rho_0 L^3}{2} \bar{\theta}^2 + \text{constant}. \quad (27)$$

Elementary quantum mechanics then tells us that for $|\bar{\theta}| \ll \pi$, the ground state wavefunction is

$$\Psi_{\Omega}(\bar{\theta}) \propto \exp(-(J\rho_0)^{1/2} L^3 \bar{\theta}^2 / 2\hbar) \quad (28)$$

which clearly becomes increasingly narrow as L increases, tending to a delta-function as $L \rightarrow \infty$.

The symmetry-breaking term also gives a way of making sense of effective-action methods. The effective-action formalism is in fact slightly more general than I described in items (1) to (3) in section 1: solving the equation

$$\left. \frac{\delta\Gamma[\phi; T]}{\delta\phi} \right|_{\phi=\phi_c} = J \quad (29)$$

gives the vacuum expectation value of the field given external field J as described by the Lagrangian (25), and the second variational derivative of \mathcal{L} evaluated at this solution gives the particle spectrum given that external field. (So the results of section 1 are the special case where $J = 0$.) For sufficiently large J , the expectation value can be nonzero, i. e. the symmetry is explicitly broken. If we define an effective action per unit volume,

$$\tilde{\Gamma} = \frac{1}{L^3} \Gamma, \quad (30)$$

then we can take L and J to zero simultaneously while keeping $J\phi_0 \gg e^{-SL^3/\hbar}$, and get a well-defined effective action in the joint limit of exact symmetry and infinite volume. It is at least plausible to suppose (though, in the absence of physically-realistic interacting models of the AQFT axioms, not rigorously proven) that this effective action describes the physics of the actually-infinite, AQFT models discussed in the previous section.

These mathematical results suggest a physical way of understanding symmetry breaking in finite systems. If we suppose that in realistic systems the system's environment is never exactly uniform, it becomes possible to suppose that for any specific system, the symmetry is explicitly broken by some environmental asymmetry. And since for sufficiently large systems the required asymmetry can be very small indeed, the supposition that such asymmetry exist might seem innocent — indeed, the contrary assumption of exact environmental asymmetry becomes extremely implausible.

This external-perturbation approach to symmetry breaking has been widely advocated both in textbooks and pedagogical articles (e.g., van Wezel and van den Brink 2007) and in conceptual and foundational work (e.g., Landsman 2013, Fraser 2015; see also further references in Fraser, *ibid*). But I do not think it is ultimately viable.

One potential problem is that in *high-energy-physics contexts* it is not so obvious that it makes sense to assume any “external” field in this sense. If we suppose that symmetry breaking occurs dynamically in the early Universe, normal physics practice is to treat that symmetry breaking as a *source* of subsequently observed inhomogeneities in the Universe, not to simply posit those inhomogeneities outright. And (unlike in condensed-matter contexts) we cannot straightforwardly regard J as representing some unanalysed external perturbation, because there is nothing obviously external to do the perturbing.

However, I do not think this is fatal to the external-perturbation approach, for a technical reason: there do not appear to be any spontaneously-broken exact symmetries in high-energy physics, at least in the sense considered in this article. (The most prominent application of global SSB in high-energy physics is the identification of pions as pseudo-Goldstone bosons associated to a spontaneous breaking of chiral symmetry in QCD — cf Weinberg (1995b, ch.19.4-19.5), Scherer (2003), Donoghue, Golowich, and Holstein (2014, ch.VII) and references therein — but there the symmetry is *explicitly* broken by the masses of the quarks, and that explicit symmetry-breaking term will certainly suffice to induce a unique noninvariant ground state.) Indeed, it is commonly said in high-energy physics (cf, e.g., Duncan (2012, pp.489-492)) that there *are* no exact global symmetries: apparent global symmetries are “accidental” symmetries that arise from the fact that symmetry-breaking terms would have to be nonrenormalisable and so are too small to be readily detected at energies accessible to us. High-energy physics does involve spontaneous breaking of exact *local* symmetries,⁷ and cosmology seems to involve spontaneous breaking of translation and rotation symmetry in *non-equilibrium* contexts⁸ but neither phenomenon lies within the scope of those discussed here.

A much more serious problem for the external-perturbation approach comes from the more mundane context of condensed-matter physics. Consider, for definiteness, a ferromagnet like an iron bar: above its critical temperature, the spins of the iron atoms are randomly allocated, but when it is cooled below the critical temperature then nearby spins line up: that is, rotational symmetry is spontaneously broken. The external-perturbation approach, taken completely literally, assumes the existence of some small constant J (a magnetic field, in effect) that explicitly breaks the symmetry, so that all of the spins in the iron bar line up in the direction of J . But this is not what happens when iron is actually cooled (unless experimenters explicitly choose to introduce a significant magnetic field). In practice, the iron bar forms “domains”, typically $\sim 10^4 - 10^6$

⁷Specifically, through the Higgs mechanism for electroweak symmetry breaking: see, e.g., Weinberg (1995b, ch.21.3)

⁸For technical details see, e.g., Weinberg (2008) and references therein; for some consideration from a philosophical point of view, see Wallace (2016a).

atoms across; atoms within a given domain are aligned with one another, but the various domains are independently and randomly aligned.

Of course, it is fairly straightforward to imagine how to modify the external-perturbation theory to account for this. We could suppose that (i) the external perturbation varies randomly from point to point, and (ii) it is constant over scales of order the domain size, so that within each domain the external perturbation determines the spin direction. This is now a quite substantial hypothesis about ambient random magnetic fields; it also implies that cooling ferromagnets are extraordinarily sensitive detectors of those random magnetic fields. Were we to take the hypothesis seriously, we should be allowing ferromagnets to cool, and measuring their domain sizes, in various different conditions where no macroscopically-detectable magnetic field is present, to see if we can detect overall anisotropies, and variations in fluctuation scale, in this mysterious background field.

But in fact, domain size does not vary in this way: it is an intrinsic feature of ferromagnets, and indeed one that is pretty well understood theoretically. Domain theory (see, e. g. , Kittel and Galt (1956) and references therein) explicitly derives the size of ferromagnetic domains as a consequence of the trade-off between the energetic costs of domain wall formation and of dipole-dipole interaction between domains. Absent a miraculous coincidence whereby space is filled with a tiny randomly-fluctuating magnetic field whose correlation length just so happens to accurately match the exact correlation length independently determined by domain theory, the account of SSB and domain formation in the ferromagnet seems to explicitly require the *absence* of any magnetic field strong enough to explicitly break symmetry and overcome the internal features of the ferromagnet that force domains to be randomly aligned on an intrinsically-determined scale.

So, while the introduction of an external field provides a satisfactory *mathematical* account of the large-volume limit of systems with SSB, it does not seem possible to interpret it as the *physical* mechanism, at least not in some important physical cases.

5 Symmetry breaking as a correlation phenomenon: the classical case

To see what I think is the right way to understand SSB in finite systems, it will help to think about the probability distribution of properties of those systems — and, for the moment, to ignore quantum mechanics and treat those systems as classical. So consider a classical model of a ferromagnet — where the system is a finite lattice of spins, each of which can point (for simplicity) either up or down — and suppose that we measure the value of *one* of those spins.

If the system's state does not spontaneously break symmetry, then of course we expect the spin to be up with probability 50%, down with probability 50%. What if it *does* break symmetry? Then *either* all the spins (or all the spins

in some domain, at any rate) are up, *or* all of them are down. But given that the overall dynamics is symmetry preserving, and given *ex hypothesi* that no external perturbation has explicitly broken the symmetry, the probability of each symmetry-breaking configuration should be equal.⁹ So *whether or not* the symmetry is broken, we predict the same statistics for each spin. And the argument generalises: given any property of the system, we expect the probability distribution over values of that property to be invariant under the symmetry, whether or not symmetry is spontaneously broken.

So where does the symmetry breaking show up? In these places:

1. In correlations between spins at different spatial points. In unbroken symmetry, the spin of one particle carries no information about the spin of other particles: spin values at different points in the lattice are uncorrelated. In spontaneously broken symmetry (and in the idealisation where domains are infinite), the spin of one particle completely determines — i. e., is perfectly correlated with — the spin of all the other particles.
2. In the symmetry-breaking properties of each individually-likely value of some collective properties, as a result of the above. The spatially-averaged spin of a large region of the ferromagnet, if the symmetry is unbroken, will with probability ~ 1 be very close to zero. The spatially-averaged spin of that same region, given spontaneously broken symmetry, will with probability ~ 1 be very close either to $+1$ or to -1 . (Note that whether or not the symmetry is spontaneously broken, the *expected* value of the spatially-averaged spin is exactly zero.)
3. In correlations between spins of the same point at different times. In unbroken symmetry, we expect the spin of any given point to be frequently re-randomised by interactions with its neighbors, so that spins at time t and time t' are uncorrelated provided that $(t - t')$ exceeds some rerandomisation timescale that is normally very short compared to other dynamically-relevant timescales. In spontaneously broken symmetry, we expect the spins at t and t' to be perfectly correlated even for very large values of $(t - t')$. (And, as a corollary, we expect the values of collective properties, like spatially-averaged spins, to be likewise almost-perfectly correlated over long times.)

It is straightforward to see explicitly how this is realised in classical statistical mechanics. Consider, to begin with, the simple case of one particle in a one-dimensional potential well described by a potential V symmetric under $x \rightarrow -x$. If that particle is in thermal contact with a reservoir, statistical mechanics describes it by the canonical distribution: the probability, at any given time, of the particle having position q and momentum p is

$$\Pr(q, p) = \frac{1}{Z(T)} e^{-p^2/2mk_B T} e^{-V(q)/k_B T}. \quad (31)$$

⁹This is admittedly a little handwaving; I will argue for it more carefully shortly.

This distribution is clearly symmetric in the spatial variable q : the expected value of particle position is zero, whether V describes a single-well potential with minimum at 0 or a double-well potential with minima at $\pm x_0$.

However, if in the latter case the height $V(0) - V(x_0)$ of the potential barrier between the two minima is large compared to $k_B T$, then the multi-time correlations in the two cases will be radically different. Because if so, in the double-well case a particle localised on the left-hand side has probability ~ 1 of having a momentum too small to cross the potential barrier and end up on the right hand side at any later time. So if we define the quantity $X(t)$, which takes value $+1$ if the particle is on the right-hand side of the system at time t (i. e., if $q(t) > 0$) and -1 if it is on the left-hand side of the system ($q(t) < 0$), we predict that $\langle X(t) \rangle = 0$ in both cases, but that $\langle X(t)X(t') \rangle$ quickly tends to zero in the single-well case, but is equal to ~ 1 even for large $(t' - t)$ in the double-well case.

As the temperature increases, of course, this will cease to be the case. If the potential barrier energy is small compared to $k_B T$ then with high probability, the particle will have sufficient energy to pass over the potential barrier with little delay, and so the correlation will fall off rapidly with time whether the potential has a single or a double minimum.

In effect, we can think of the canonical distribution of the system, at sufficiently low temperatures, as being written as

$$\Pr(q, p) = 0.5 \times \Pr_L(q, p) + 0.5 \times \Pr_R(q, p) \quad (32)$$

where \Pr_L and \Pr_R are, respectively, the appropriate canonical ensembles for a particle trapped on the left or the right of an infinite potential barrier at $x = 0$. Because there is no such barrier, these separate distributions are not *quite* time-invariant, but to a very good approximation we can consider the system as being in one of two thermal states, with probability 0.5 of each. Using the language of measurement for a moment, if we observe this system then our first observation will (among other things) determine which of the two thermal states the system is in; all subsequent measurements will be relative to that state. But all of this is captured by the single, symmetric, canonical ensemble of the whole system, and need not be put in by hand.

The situation is conceptually still clearer, if technically more complicated, if we consider a classical version of the field-theory case (5) with a double-well potential. (Since this theory needs to be discretised at high momenta in order to have a well-defined statistical mechanics, it is closely related to the classical ferromagnet we considered above.¹⁰) At absolute zero, the canonical distribution for this system is a sum of delta functions: the system has state $\phi(x) = +\phi_0$ with probability 0.5 and $\phi(x) = -\phi_0$ with probability 0.5. So the expected value of $\phi(x)$ for any x is zero, but there is perfect correlation between $\phi(x, t)$ and $\phi(x', t')$ for any x, x', t, t' . At temperatures above zero, but still sufficiently small, the canonical distribution of the system will assign probability ~ 0.5 to

¹⁰For a more precise description of the relation between them see, e. g., Binney et al (1992, chs.6-7).

the system's state being a small fluctuation around $\phi(x) = +\phi_0$, and probability ~ 0.5 to it being a similar fluctuation around $\phi(x) = -\phi_0$. And the probability (again, as determined by the canonical distribution) of a given state transitioning from the vicinity of one minimum to the vicinity of another in any reasonable timescale will be negligible. Again, we can think of the canonical ensemble of the system as an equally-weighted sum of two canonical ensembles, each *very nearly* time-invariant and each describing a system localised around one or other minimum. In any observational context, the first observation serves to determine which minimum is realised; given the negligible probability of the system transitioning between minima in any reasonable timescale, subsequent observations are determined by the thermal state relevant to the actual minimum.

In this macroscopic system (as opposed to the purely microscopic case of the single particle) we can also use the language of Boltzmannian statistical mechanics (cf, e.g., Goldstein (2001), Lebowitz (2007)) to describe what is going on. In that language, we partition the phase space of a system (more precisely, each energy hypersurface of that phase space) into 'macrostates', each characterised by approximately-constant values of collective physical quantities like (in this case) spatially-averaged field strength over some macroscopic region. On the normal Boltzmannian account of equilibration, for any given energy one such macrostate (the "equilibrium macrostate") will have a vastly larger volume than the rest, and so in the absence of other conserved quantities it is reasonable to expect that the system will make its way into that macrostate and remain there for a very long time. The validity of this assumption about the equilibrium macrostate can be checked in the canonical formalism by looking at the variance of the spatially averaged quantities according to the canonical distribution. If it is sufficiently small, the probability of these macroquantities being close to their most-probable value is very large, and — since the canonical distribution is uniform on each energy hypersurface — this at least loosely speaking corresponds to the equilibrium macrostate as being dominant in volume.

For this very reason, the straightforward Boltzmannian hypothesis — i.e., that there is a unique equilibrium macrostate dominating the phase space — will fail for this system. What we can expect instead (for sufficiently low energies) is that (i) the overwhelming majority of the volume of the hypersurface is split between two equal-volume macrostates, one corresponding to each minimum; (ii) states initially localised in one of these macrostates almost all (as measured by phase space volume) avoid crossing over to the other macrostate unless we wait an extremely long time. In this case, the Boltzmannian account of equilibration will predict that the system will end up in *one or other* of the two equilibrium macrostates. If the probability distribution over non-equilibrium initial conditions is itself invariant under the symmetry, each equilibrium has a 50% chance of being realised.

To summarise: at least in the classical case, there is a perfectly clear account of SSB to be given that is compatible both with finite systems, and with probability distributions that are required to be invariant under the symmetry.

6 Correlation probabilities for a quantum particle

In the light of these observations, let us revisit the ground state of the *quantum* systems we first considered in Section 1. In the first instance, consider the one-dimensional particle in a symmetric potential well. On symmetry grounds it is clear that the position of the particle is equally likely to be measured at $x > 0$ or $x < 0$, irrespective of whether we have a single-well or double-well potential. But investigating the multi-time correlation functions for the system is more difficult, since quantum states cannot simply be treated as probability distributions. For the moment, let us adopt an operational approach and suppose that we measure, at successive times $0, \tau, 2\tau, \dots$, whether the particle is on the left or right side of the system.

If we begin with the single-well harmonic oscillator, we can avoid technical complexities by confining our attention to the two lowest-lying energy eigenstates:

$$\psi_0(x) = \frac{1}{\pi^{1/4}} e^{-m\omega x^2/2} \quad (33)$$

and

$$\psi_1(x) = \frac{(2m\omega)^{1/2}}{\pi^{1/4}} x e^{-m\omega x^2/2}. \quad (34)$$

The sum and difference of these functions,

$$\psi_R(x) = \frac{1}{\sqrt{2}}(\psi_0(x) + \psi_1(x)) \quad \psi_L(x) = \frac{1}{\sqrt{2}}(\psi_0(x) - \psi_1(x)) \quad (35)$$

are fairly well-localised on the right and left sides of $x = 0$ respectively: if the system's state is ψ_R , the probability of it being found with $x > 0$ is $(1/2 + (2\pi)^{-1/2}) \simeq 0.899$. So to a reasonable first approximation we can replace measurement of whether the particle is on the left or right with measurement in the ψ_L, ψ_R basis.

Next, let us model a measurement in this basis by the following unitary interaction (adopting Dirac notation):

$$|\psi_L\rangle \otimes |-\rangle \rightarrow |\psi_L\rangle \otimes |L\rangle \quad |\psi_R\rangle \otimes |-\rangle \rightarrow |\psi_R\rangle \otimes |R\rangle. \quad (36)$$

Here, $|-\rangle, |L\rangle$ and $|R\rangle$ represent, respectively, states of some external measurement device that is in the “ready” state, that has recorded the particle to be in state ψ_L , and that has recorded the particle to be in state ψ_R , and the overall interaction is assumed to be so quick (relative to the dynamical timescales of the oscillator itself, i. e. ω^{-1}) that we can idealize it as instantaneous. Introducing two such measurement devices and allowing the system to evolve unitarily between the first measurement at $t = 0$ and the second at $t = \tau$, we find the state of system+devices after that second measurement to be

$$\begin{aligned} |\psi_2\rangle &= \frac{1}{\sqrt{2}} \cos(\omega\tau) (|\psi_L\rangle \otimes |L\rangle \otimes |L\rangle + |\psi_R\rangle \otimes |R\rangle \otimes |R\rangle) \\ &+ \frac{1}{\sqrt{2}} \sin(\omega\tau) (|\psi_L\rangle \otimes |R\rangle \otimes |L\rangle + |\psi_R\rangle \otimes |L\rangle \otimes |R\rangle). \end{aligned} \quad (37)$$

To extend to three measurements, we can define a condensed notation where $|XYZ\rangle \equiv |\psi_Z\rangle \otimes |X\rangle \otimes |Y\rangle \otimes |Z\rangle$. Then the quantum state after three measurements is

$$\begin{aligned} |\psi_3\rangle &= \frac{1}{\sqrt{2}} \cos^2(\omega\tau)(|LLL\rangle + |RRR\rangle) + \frac{1}{\sqrt{2}} \sin^2(\omega\tau)(|LRL\rangle + |RLR\rangle) \\ &+ \frac{1}{2} \cos(\omega\tau) \sin(\omega\tau)(|LLR\rangle + |RRL\rangle + |LRR\rangle + |RLL\rangle). \end{aligned} \quad (38)$$

Comparison of these two expressions reveals that the “probabilities” that they define do not behave like normal probabilities. Consider the following protocols: (I) measure at time 0, make no measurement at time τ , measure at time 2τ . (II) measure at each of times 0, τ , 2τ . The probability of obtaining (say) the result “left, left” in protocol (I) is (applying the Born rule to (37), with $\tau \rightarrow 2\tau$)

$$\Pr(L?L) = \frac{1}{2} \cos^2(2\omega\tau). \quad (39)$$

The probabilities of obtaining “left,right,left” and “left,left,left” in protocol (II) (applying the Born rule to (38)) are

$$\begin{aligned} \Pr(LLL) &= \frac{1}{2} \cos^4(\omega\tau) = \frac{1}{8} (1 + \cos(2\omega\tau))^2 \\ \Pr(LRL) &= \frac{1}{2} \sin^4(\omega\tau) = \frac{1}{8} (1 - \cos(2\omega\tau))^2 \end{aligned} \quad (40)$$

so that

$$\Pr(LLL) + \Pr(LRL) = \frac{1}{4} (1 + \cos^2(2\omega\tau)) \quad (41)$$

which in general does not equal $\Pr(L?L)$.

This, of course, is nothing but quantum interference: there is interference between the LLL and LRL histories such that the sum of their individual probabilities is not necessarily the probability of the more coarse-grained history. When interference is present, it is not mathematically possible to interpret the probabilities defined by our operational protocol as the probabilities that the unmeasured system actually has the measured values of the quantities being measured.

If we apply a similar analysis to the double-well potential, *formally* we get the same answer. We can define ψ_L and ψ_R in terms of the ground state and the first excited state (as we have already seen, these states are very well approximated, for sufficiently widely separated wells, by the ground states of simple harmonic oscillators centred at the minima of the two wells); the formulae just given will again apply to correlation measurements; the only difference is that instead of ω , we use the frequency gap

$$\omega_I = \kappa\omega(S/\hbar)^{1/2} e^{-S/\hbar} \quad (42)$$

that (as we discussed in Section 2) can be calculated by instanton methods. However, there is a crucial *substantive* difference. If the wells are sufficiently

separated that $S/\hbar \gg 1$, then $\omega_I \ll \omega$. It is then possible to choose τ such that $\omega^{-1} \sim \tau \ll \omega_I^{-1}$: that is, the time between measurements is comparable to the oscillation timescales of the individual wells but much less than $1/\omega_I$, the tunnelling timescale for transition between wells. In this circumstance, the probability of obtaining different results on different measurements becomes negligibly small: if the system is measured and found to be on the left hand side, it will continue to be found on the left hand side in any future measurements. As such there is no appreciable interference, and thus no formal barrier — at least as far as these measurements are concerned — to interpreting the ground state of the double-well oscillator as a probabilistic mixture of “L” and “R”.

This is suggestive: just as in the classical case, there seems to be a difference between single-well and double-well potentials that manifests in high correlation between particle locations at different times. Indeed, the difference between the two cases seems to manifest not only in different values of the correlations between outcomes, but in the formal validity or invalidity of attributions of probabilities to sequences of outcomes in the first place. In effect, in the single-well case, if we want to study the system on the timescales on which its dynamics is non-trivial, we have to treat its quantum state as a description of the actual, and highly non-classical, physical state. In the double-well case, provided we are not interested in timescales long enough for tunnelling to matter, we are at least formally justified in treating the state as a probabilistic mixture of two single-well-potential quantum states, just as in the classical case.

In the rest of the paper I will argue that these observations can be developed into a fully satisfactory understanding of SSB in finite systems.

7 Decoherence and consistent histories: a brief review

The questions of whether, when, why, and to what degree the quantum state can be treated as a probability distribution is closely tied to the quantum measurement problem: indeed (at least in my view; see Wallace (2016b) for more details) the most accurate statement of the measurement problem is exactly this set of questions. As such, our consideration of SSB has now taken us into deep and contentious waters.¹¹ But there is at least a well-established formalism to study when the quantum state can *safely be treated* as probabilistic: the *decoherent-histories formalism* (Griffiths (1984, 1993), Gell-Mann and Hartle (1989, 1993), Halliwell (1995), Omnes (1992)) which I will briefly describe for readers unfamiliar with it.

The ingredients for the decoherent-histories formalism (in a slightly simplified form sufficient for our purposes) are:

¹¹The idea that there might be some link between the measurement problem, inequivalent representations, and SSB is not new; it has been previously explored by, e. g., Hepp (1972), Saunders (1988), Landsman (1991), and Breuer *et al* (1993). However, these papers develop a quite distinct line of thought from the one I develop here: in particular, appeal to the infinite limit appears essential in each.

1. A quantum theory, defined on some Hilbert space \mathcal{H} and with dynamics given by a family of unitary operators $\hat{U}(t)$;
2. An increasing sequence of times t_0, t_1, \dots, t_n ;
3. An orthogonal family of projectors $\{\hat{\Pi}_J\}$, such that $\sum_J \hat{\Pi}_J = \hat{1}$, representing mutually-exclusive properties of the system;
4. An initial state — that is, a density operator ρ on \mathcal{H} — representing the system's state at time t_0 .

A *history* h of the system is then a function from times t_i in the sequence to operators $h(t_i)$, where each $h(t_i)$ is a sum of some subset of the projectors $\hat{\Pi}_J$, and hence represents a possibly-coarse-grained property. The intended interpretation is that $h(t_i)$ is the property possessed by the system at time t_i . The histories are then exactly the sequences of (possibly-coarse-grained) ideal measurements (that is, measurements of the kind we discussed above) of the $\hat{\Pi}_J$. To each history corresponds a *history operator*

$$\hat{H}(h) = h(t_n)\hat{U}(t_n - t_{n-1})h(t_{n-1}) \cdots \hat{U}(t_2 - t_1)h(t_1)\hat{U}(t_1 - t_0)h(t_0). \quad (43)$$

We now define the *decoherence functional* as

$$D_\rho(h, h') = \text{Tr}(\rho \hat{H}^\dagger(h) \hat{H}(h')). \quad (44)$$

(Note that the functional depends on the initial state as well as on the histories.)

The decoherence functional has two key properties. Firstly, $D_\rho(h, h) \equiv \text{Pr}_\rho(h)$ is the Born-rule probability that a sequence of ideal measurements of the properties represented by the $\hat{\Pi}_J$ will give a sequence of outcomes of which history h is a coarse-graining. Secondly, suppose that h_1 and h_2 satisfy $h_1(t) = h_2(t)$ except for $t = t_i$, and that $h_1(t_i)$ and $h_2(t_i)$ are orthogonal. Then we can define a history $h_1 + h_2$ by

$$\begin{aligned} h(t) &= h_1(t) = h_2(t) \quad t \neq t_i \\ h(t_i) &= h_1(t_i) + h_2(t_i). \end{aligned} \quad (45)$$

Classically, we would expect $\text{Pr}_\rho(h_1 + h_2) = \text{Pr}_\rho(h_1) + \text{Pr}_\rho(h_2)$, but quantum-mechanically we find

$$\begin{aligned} \text{Pr}_\rho(h_1 + h_2) &= \text{Pr}_\rho(h_1) + \text{Pr}_\rho(h_2) + D_\rho(h_1, h_2) + D_\rho(h_2, h_1) \\ &= \text{Pr}_\rho(h_1) + \text{Pr}_\rho(h_2) + 2\text{Re } D_\rho(h_1, h_2). \end{aligned} \quad (46)$$

So the real part of the decoherence functional between *different* histories is a measure of the interference between those histories. If $\text{Re } D_\rho(h, h') = 0$ for distinct histories h, h' , we say that the space of histories is *consistent* — meaning that we can consistently use the decoherence function to assign probabilities to those histories (at least formally). In practice, the weaker criterion of *approximate consistency* — $\text{Re } D_\rho(h, h') \simeq 0$ for $h \neq h'$ — is generally taken to be

sufficient, guaranteeing as it does that $D(h, h)$ can for all practical purposes be treated as a probability. A more calculationally useful, though strictly speaking overly strong, variation on this is *medium decoherence*, the requirement that $D_\rho(h, h') \simeq 0$.

Of course, identifying formal conditions under which the quantum state can be treated as probabilistic does not in itself resolve the deep conceptual puzzle about how to understand a description of a system that is probabilistic in some contexts and seems to represent the objective physical state of the system in other contexts. (My own view, defended at length in Wallace (2012), is that we need to understand the quantum state as representing the objective, physical state of the system in all contexts, and to see the medium decoherence condition as a necessary condition to interpret that state as realising a branching, emergently probabilistic structure along the lines of Everett's (1957) many-worlds theory — but I shall not presuppose that view here.) But however this conceptual problem is resolved, to make contact with the empirical results of quantum mechanics it will have to — somehow — legitimate the normal practice of interpreting the quantum state as probabilistic in the absence of interference. So I will assume, for the purposes of this article, that we can interpret decoherent quantum histories probabilistically.

So much for the *definition* of decoherence and consistency. The most commonly discussed *dynamical origin* of consistency is *environment-induced decoherence* (Zeh (1973, 1993), Joos and Zeh (1985), Zurek (1991, 1998); for reviews, see Joos *et al* (2003) or Schlosshauer (2007)), where the family of projectors tracks only properties of some system and traces over the properties of an environment (this was proposed as a mechanism for SSB by Simonius (1978)).

But (as stressed by, e.g., Halliwell (1998, 2010)) another source of decoherence is *conservation*: if some quantity is approximately conserved by the dynamics, the histories of that quantity will decohere.

In more detail: suppose that we choose our family of projectors $\{\hat{\Pi}_J\}$ to satisfy $\hat{U}(t)\hat{\Pi}_J\hat{U}^\dagger(t) \simeq \hat{\Pi}_J$. Then we can distinguish between constant histories (for which $h(t) = h(t')$) and non-constant histories. It is easy to check that the decoherence functional now simplifies dramatically, to

$$\begin{aligned} D_\rho(h, h') &\simeq 0 \quad (h \text{ or } h' \text{ non-constant or } h \neq h') \\ &\simeq \text{Tr}(\rho h(t_0)) \quad (\text{constant } h, h = h') \end{aligned} \tag{47}$$

In this situation, we can treat the system as decomposing into dynamically isolated sectors, each characterised by a particular value of the quantity represented by the $\{\hat{\Pi}_i\}$, and each occurring with probability $\text{Tr}(\rho\hat{\Pi}_i)$. As we will shortly see, this conservation-based conception of decoherence plays a central role in SSB.

8 Decoherence and spontaneous symmetry breaking

Let's consider the single- and double-well potentials from a decoherent-histories viewpoint. We choose as our family of projectors the two projectors $\hat{\Pi}_L, \hat{\Pi}_R$ that project onto states localised on (respectively) the left and right sides of the origin, choose as times the equally-spaced times $0, \tau, 2\tau, \dots, n\tau$, and continue to work in the approximation introduced in Section 6 where we consider only the ground and first excited states $|\psi_0\rangle, |\psi_1\rangle$ (with energy separation $\hbar\omega$), so that we can approximate

$$\hat{\Pi}_L = \frac{1}{2}(|\psi_0\rangle - |\psi_1\rangle)(\langle\psi_0| - \langle\psi_1|)$$

$$\hat{\Pi}_R = \frac{1}{2}(|\psi_0\rangle + |\psi_1\rangle)(\langle\psi_0| + \langle\psi_1|).$$

Given the results of Section 6, if $t \ll 1/\omega$ then $\hat{U}(t)\hat{\Pi}_{(L/R)}\hat{U}^\dagger(t) \simeq \hat{\Pi}_{(L/R)}$, so that the conservation conditions for consistency are satisfied. We can then reinterpret section 6 as telling us that for the double-well potential, unless we study it on timescales long compared to the between-well tunnelling timescale, the left-side and right-side histories decohere and the system can thus be considered as a probabilistic mixture of a particle localised on the left side and one localised on the right side. Furthermore, this should continue to hold if we replace the ground state with a thermal state with expected energy much lower than the height of the energy barrier between the two wells (since such a state is a thermal mixture of energy eigenstates that come in symmetric/antisymmetric pairs with extremely small energy gaps between them, just as for the ground state).

It is important to note that the system is by no means *completely* decoherent, nor is its full dynamics even approximately classical. The projectors $\hat{\Pi}_L, \hat{\Pi}_R$ are extremely coarse-grained, and the left-side and right-side sectors can still display complicated quantum-mechanical dynamics (the dynamics of a perturbed harmonic oscillator, in particular). The decoherence between left-side and right-side histories allows the quantum state to be treated as a probabilistic mixture of left-side and right-side quantum states, but those states are not further decohered, at least by any process internal to the oscillator.

As in the classical case, things become conceptually clearer (if technically more complicated) if we move to the case of a field theory with a double-well potential. The ground state of that system is a classic ‘‘Schrödinger-cat’’ state: a superposition of two wavefunctions localised around macroscopically distinct field configurations. As such, even before considering its dynamics we can see that the system displays the first two (closely related) characteristics of SSB identified in Section 5: the field values of different spatial points will be closely correlated, and the spatially averaged field strength, though it has expected value zero, will with probability ~ 1 be found to be macroscopically different from zero.

As for the dynamical behavior of the system: formally speaking, the analysis we have already made applies here too, simply substituting for the energy gap ΔE the result (18) obtained in section 2. Provided that we are considering the dynamics on timescales $t \ll \hbar/\Delta E$, and provided the system's initial state has expected energy well below the energy required for thermal transitions across the potential barrier, the system will decohere into two sectors characterised by positive and negative average field values. Each sector will look like a scalar quantum field theory describing excitations around the respective minima, and dynamical coupling between the two sectors will be so negligible that interference between them may be neglected, justifying the probabilistic interpretation. The two sectors individually, however, will be fully quantum-mechanical (at least in the absence of some further decoherence mechanism).

It is instructive to quantify the timescales on which this treatment is appropriate, for a solid-state system. They are controlled by the instanton action S defined by equation (19), according to which the action is proportional to the volume of the system. We can obtain an extremely crude estimate simply via dimensional analysis: solid-state systems are characterised by typical masses m around the proton mass m_p , typical energies E around the 1eV typical of atomic systems, and typical lengthscales l around the 10^{-10} m interatomic distance. The most general way to combine these to get a quantity of dimension (action/volume) is

$$S/V \alpha \frac{\sqrt{mE}}{l^2} \sim \alpha \times 10^{-3} \text{Js/m}^3 \quad (48)$$

in SI units, where α is a dimensionless constant that (in the absence of really remarkable emergent effects) will be within a few orders of magnitude of unity. (It will be apparent in a moment why this level of crudeness is acceptable). The overall dependence of ΔE on S/V is, to leading order and ignoring dimensionless factors,

$$\Delta E \sim E \times e^{-SL^3/\hbar V}. \quad (49)$$

For (say) a one-cubic-millimeter sample of a solid, $SL^3/\hbar V \sim 10^{22}$, so that

$$\Delta E \sim e^{-10^{22}}. \quad (50)$$

We would then expect the two symmetry-breaking histories to be effectively decoherent for timescales short compared to $e^{10^{22}}$ s. To call this length of time “large compared to the age of the Universe” would be a ridiculous understatement. So even for a finite — indeed, quite small — sample, and even in the complete absence of external perturbations, we can to a ridiculously good approximation treat a system below the phase-transition energy as being in a probabilistic mixture of two symmetry-breaking states.

The decoherent-histories framework can also be applied to cases of spontaneously-broken *continuous* symmetry, although there is no some arbitrariness in the definition of the history projectors. Starting from the nonlinear-sigma-model approximation in equation (23), we need to define projectors with respect to the spatially-averaged phase θ , and we can do so by dividing the unit circle

into equal sections of angular width Δ and defining projectors onto states with spatially-averaged phase having support in those sections. In this case (see Appendix for details), the conservation condition for consistent histories is satisfied on timescales t satisfying

$$t \ll \frac{\Delta^2 L^3 \rho_0^2}{\hbar}. \quad (51)$$

This goes to infinity as $L \rightarrow \infty$; indeed, we can take a simultaneous limit of $\Delta \rightarrow 0$ and $L \rightarrow \infty$, and recover the fact that in the infinite limit the dynamically-isolated sectors are identified with a sharp value of the spatially-averaged phase. Furthermore, although it does not go to infinity nearly so quickly as in the discrete case (it diverges linearly, rather than exponentially, with system volume) there are good reasons to think that it is sufficiently quick to explain the actual phenomenology of finite systems. Using the same crude dimensional-analysis approach used previously, we find that ρ_0^2 has order of magnitude

$$\rho_0^2 \sim \frac{\hbar}{l} \left(\frac{m}{E}\right)^{1/2} \quad (52)$$

so that the timescale of validity for the consistent-history approach is

$$t_{CH} \sim l \left(\frac{m}{E}\right)^{1/2} \left(\frac{L}{l_0}\right)^3 \Delta^2 \sim 10^{16} L^3 \Delta^2 \quad (53)$$

in SI units. As a rough estimate for the timescales t_μ relevant for microphysical processes in a solid, we can take $t_\mu = L/c$ where $c \sim 10^3$ m/s is the speed of sound in the solid. So we can treat the histories as decohering when studying solids provided that $t_\mu \ll t_{CH}$, i. e.

$$\Delta^2 \gg 10^{-19}/L^2. \quad (54)$$

For the one-cubic-millimeter sample we previously considered, this becomes

$$\Delta^2 \gg 10^{-11} \quad (55)$$

as a condition for applying the decoherent-histories formalism. SSB will occur provided $\Delta \ll 2\pi$, so (on this admittedly crude estimate) we can treat the sample as spontaneously breaking symmetry on timescales far longer than the relevant microphysical timescales for the system, even for a small sample.

For a sufficiently hot sample, and or a sufficiently energetic external environment, we would also expect environment-induced decoherence to induce SSB, because either thermally-excited Goldstone bosons or the external environment will record the spatially-averaged phase. In principle, an extremely small sample (only a few dozen atoms across) cooled to a sufficiently low temperature and sufficiently isolated from its environment might not be well treated by assuming broken symmetry, but systems like this in any case lie outside the class of systems that we would expect to be well treated by conventional field-theoretic methods.

9 Conclusion

According to Butterfield’s account, we can only understand SSB via the infinite limit if we can identify some feature of finite systems which corresponds at least approximately to SSB, with the approximation becoming exact in the limit. Degeneracy of the ground state is not such a feature, appearing as it does only for actually-infinite systems. Explicit symmetry breaking by external perturbations has the right formal structure, but seems incompatible with features of the actual phenomenology of symmetry breaking in finite systems.

The real signature of SSB in finite systems is the breakdown of those systems’ state spaces, at least at low energies, into sectors that are approximately dynamically isolated, as identified using the tools of the decoherent-histories approach. The approximation becomes exact in the infinite limit: we can interpret the inequivalent representations that correspond to spontaneous symmetry breaking as infinite-volume limits of the approximately-isolated decoherent sectors. (In that limit, different-temperature regimes also correspond to inequivalent representations, so the restriction to low energy ceases to apply: the symmetry is still restored at high temperature, but this is seen at the level of representations rather than within any given representation.)

From this perspective, the common identification of SSB with ground-state degeneracy looks misleading. Ground state degeneracy matters because, in an infinite system, it is roughly-speaking equivalent to the existence of inequivalent, symmetry-related vacuum representations, but it is the existence of those representations, not ground-state degeneracy per se, which gives rise to the phenomenology generally associated to spontaneous symmetry breaking.

Also from this perspective, we can understand the close relation between SSB and the quantum measurement problem, even without appeal to the infinite limit. A quantum system, initially in a perfectly symmetric quantum state describing thermal equilibrium above the critical temperature, remains in that perfectly symmetric state as it is cooled to absolute zero. What changes is that it realises that symmetry, at sufficiently low temperatures, by being a superposition of macroscopically distinct states, each individually non-symmetric, and each acting as the *de facto* ground state of a space of states dynamically isolated from those spaces built on other terms in the superposition. Whatever our preferred resolution to the measurement problem, it will need to permit us to treat such macroscopic superpositions as probabilistic mixtures; whether that is achieved through some new dynamics or new variables, through some probabilistic reinterpretation of the quantum state, or through the emergent branching of the Everett interpretation, it will also provide us with what we need to resolve the related puzzle of spontaneous symmetry breaking in finite systems.

Acknowledgements

I’m extremely grateful to Jeremy Butterfield for a close and careful reading of the manuscript, which has improved it substantially.

Appendix

Consider a quantum-mechanical particle in one dimension with Hamiltonian

$$\hat{H} = \hat{P}^2/2m. \quad (56)$$

and initially in a square-wave state,

$$\begin{aligned} \langle x|\psi\rangle &= \frac{1}{\sqrt{\Delta}} \quad (|x| < \Delta/2) \\ &= 0 \quad (|x| > \Delta/2) \end{aligned} \quad (57)$$

In a momentum basis, this state is

$$\langle k|\psi\rangle = \sqrt{\frac{2}{\pi\Delta}} \frac{\sin(k\Delta/2)}{k}. \quad (58)$$

So

$$\langle\psi|\exp(-it\hat{H}/\hbar)|\psi\rangle = \frac{2}{\pi\Delta} \int dk e^{-itk^2/2m\hbar} \frac{\sin^2(k\Delta/2)}{k^2} \quad (59)$$

or, after a change of variable,

$$\langle\psi|\exp(-it\hat{H}/\hbar)|\psi\rangle = \frac{2}{\pi} \int dz e^{-itz^2/2m\Delta^2\hbar} \frac{\sin^2(z)}{z^2}. \quad (60)$$

If $t \ll m\Delta^2\hbar$, then the exponent in the integral will vary only slowly in the region where $\sin^2(z)/z^2$ is large, and so may be approximated as constant. In this regime, the integral is t -independent and so can be evaluated at $t = 0$, where it is trivially equal to 1. So a square wave-packet does not spread out appreciably provided that this condition is satisfied.

The spatially-averaged phase of the field theory describing spontaneous symmetry breaking can be approximated by this system if we put $m = \rho_0^2 L^3$, provided we take $\Delta \ll 2\pi$ (so that the system may be approximated as permitting phase values in $(-\infty, \infty)$ and not just $[0, 2\pi)$). We then have our condition for consistency of histories defined by projections of width Δ :

$$t \ll \frac{\Delta^2 L^3 \rho_0^2}{\hbar}. \quad (61)$$

References

- Altland, A. and B. D. Simons (2010). *Condensed Matter Field Theory* (2nd ed.). Cambridge: Cambridge University Press.
- Anderson, P. W. (1984). *Basic Notions of Condensed Matter Physics*. Westview Press.
- Ashcroft, N. W. and N. D. Mermin (1976). *Solid State Physics*. Orlando, FL.: Harcourt.

- Baker, D. J. and H. Halvorson (2013). How is spontaneous symmetry breaking possible? understanding Wigner’s theorem in light of unitary inequivalence. *Studies in the History and Philosophy of Modern Physics B* 44, 464–469.
- Batterman, R. W. (2001). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction and Emergence*. Oxford: Oxford University Press.
- Batterman, R. W. (2005). Critical phenomena and breaking drops: Infinite idealizations in physics. *Studies in History and Philosophy of Modern Physics* 36, 225–244.
- 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.
- Breuer, T., A. Amann, and N. P. Landsman (1993). Robustness in quantum measurements. *Journal of Mathematical Physics* 34, 5441.
- Butterfield, J. (2011). Less is different: Emergence and reduction reconciled. *Foundations of Physics* 41, 1065–1135.
- Coleman, S. (1985). *Aspects of Symmetry: selected Erice lectures*. Cambridge, UK: Cambridge University Press.
- Coleman, S. and E. Weinberg (1973). Radiative corrections as the origin of spontaneous symmetry breaking. *Physical Review D* 7, 1888.
- DeWitt, B. and N. Graham (Eds.) (1973). *The many-worlds interpretation of quantum mechanics*. Princeton: Princeton University Press.
- Dolan, L. and R. Jackiw (1974). Symmetry behavior at finite temperature. *Physical Review D* 9, 3320–3341.
- Donoghue, J. F., E. Golowich, and B. R. Holstein (2014). *Dynamics of the Standard Model (Second Edition)*. Cambridge: Cambridge University Press.
- Duncan, T. (2012). *The Conceptual Framework of Quantum Field Theory*. Oxford: Oxford University Press.
- Earman, J. (2004). Curie’s principle and spontaneous symmetry breaking. *International Studies in the Philosophy of Science* 18, 173–198.
- Everett, H. I. (1957). Relative State Formulation of Quantum Mechanics. *Review of Modern Physics* 29, 454–462. Reprinted in DeWitt and Graham (1973).
- Feynman, R. P. and A. R. Hibbs (1965). *Feynman, Richard P. and Hibbs, Albert R.* New York: McGraw-Hill.
- Fraser, J. (2015). Spontaneous symmetry breaking in finite systems. *Philosophy of Science* 83, 585–605.
- Gell-Mann, M. and J. B. Hartle (1989). Quantum Mechanics in the Light of Quantum Cosmology. In W. H. Zurek (Ed.), *Complexity, Entropy and the*

- Physics of Information*, pp. 425–459. Redwood City, California: Westview Press.
- Gell-Mann, M. and J. B. Hartle (1993). Classical equations for quantum systems. *Physical Review D* 47, 3345–3382.
- Goldstein, S. (2001). Boltzmann’s approach to statistical mechanics. In J. Bricmont, D. Dürr, M. Galavotti, F. Petruccione, and N. Zanghi (Eds.), *In: Chance in Physics: Foundations and Perspectives*, Berlin, pp. 39. Springer. Available online at <http://arxiv.org/abs/cond-mat/0105242>.
- Griffiths, R. (1993). Consistent interpretation of quantum mechanics using quantum trajectories. *Physical Review Letters* 70, 2201–2204.
- Griffiths, R. B. (1984). Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics* 36, 219–272.
- Haag, R. (1996). *Local Quantum Theory: Fields, Particles, Algebras*. Berlin: Springer-Verlag.
- Halliwel, J. J. (1995). A review of the decoherent histories approach to quantum mechanics. *Annals of the New York Academy of Sciences* 755, 726–740.
- Halliwel, J. J. (1998). Decoherent histories and hydrodynamic equations. *Physical Review D* 35, 105015.
- Halliwel, J. J. (2010). Macroscopic superpositions, decoherent histories and the emergence of hydrodynamic behaviour. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford: Oxford University Press.
- Halvorson, H. (2007). Algebraic quantum field theory. In J. Butterfield and J. Earman (Eds.), *Handbook of the Philosophy of Science: Philosophy of Physics, Part A*. Boston: Elsevier.
- Hepp, K. (1972). Quantum-theory of measurement and macroscopic observables. *Helvetica Physica Acta* 45, 237.
- Joos, E. and H. Zeh (1985). The emergence of classical particles through interaction with the environment. *Zeitschrift für Physik B59*, 223–243.
- Joos, E., H. D. Zeh, C. Kiefer, D. Giulini, J. Kupsch, and I. O. Stamatescu (2003). *Decoherence and the Appearance of a Classical World in Quantum Theory* (2nd Edition ed.). Berlin: Springer.
- Kittel, C. and J. Galt (1956). Ferromagnetic domain theory. *Solid State Physics* 3, 437–564.
- Landsman, K. (2013). Spontaneous symmetry breaking in quantum systems: Emergence or reduction? *Studies in the History and Philosophy of Modern Physics* 44, 379–394.
- Landsman, N. P. (1991). Algebraic theory of superselection sectors and the measurement problem in quantum mechanics. *International Journal of Modern Physics A* 6, 5349.

- Lebowitz, J. (2007). From time-symmetric microscopic dynamics to time-asymmetric macroscopic behavior: An overview. Available online at <http://arxiv.org/abs/0709.0724>.
- Omnes, R. (1992). Consistent interpretations of quantum mechanics. *Reviews of Modern Physics* 64, 339–382.
- Peskin, M. E. and D. V. Schroeder (1995). *An introduction to Quantum Field Theory*. Reading, Massachusetts: Addison-Wesley.
- Ruetsche, L. (2003). A matter of degree: Putting unitary inequivalence to work. *Philosophy of Science* 70, 1329–1342.
- Ruetsche, L. (2006). Johnny’s so long at the ferromagnet. *Philosophy of Science* 73, 473–486.
- Ruetsche, L. (2011). *Interpreting Quantum Theories*. Oxford: Oxford University Press.
- Saunders, S. (1988). The algebraic approach to quantum field theory. In H. Brown and R. Harre (Eds.), *Philosophical Foundations of Quantum Field Theory*. Oxford: Clarendon Press.
- Scherer, S. (2003). Introduction to chiral perturbation theory. *Advances in Nuclear Physics* 27, 277–538.
- Schlosshauer, M. (2007). *Decoherence and the Quantum-to-Classical Transition*. Berlin: Springer.
- Simonius, M. (1978). Spontaneous symmetry breaking and blocking of metastable states. *Physical Review Letters* 40, 980–983.
- van Wezel, J. and J. van den Brink (2007). Spontaneous symmetry breaking in quantum mechanics. *American Journal of Physics* 75, 635.
- Wallace, D. (2011). Taking particle physics seriously: a critique of the algebraic approach to quantum field theory. *Studies in the History and Philosophy of Modern Physics* 42, 116–125.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.
- Wallace, D. (2016a). Interpreting the quantum mechanics of cosmology. Forthcoming in “Guide to the Philosophy of Cosmology”, Ijjas and Loewer (ed.).
- Wallace, D. (2016b). What is orthodox quantum mechanics? available online at <https://arxiv.org/abs/1604.05973>.
- Wallace, D. (2017). Quantum field theory. Forthcoming in Knox and Wilson (ed.), *Handbook of the Philosophy of Physics* (Routledge).
- Weinberg, S. (1973). Perturbative calculations of symmetry breaking. *Physical Review D* 7, 2887.
- Weinberg, S. (1995a). *The Quantum Theory of Fields, Volume I: Foundations*. Cambridge: Cambridge University Press.

- Weinberg, S. (1995b). *The Quantum Theory of Fields, Volume II: Modern Applications*. Cambridge: Cambridge University Press.
- Weinberg, S. (2008). *Cosmology*. Oxford: Oxford University Press.
- Witten, E. (2000). Dynamics of quantum field theory. In P. Deligne, P. Etingof, D. S. Freed, L. C. Jeffrey, D. Kazhdan, J. W. Morgan, D. R. Morrison, and E. Witten (Eds.), *Quantum Fields and Strings: a Course for Mathematicians*, Volume 2, pp. 1119–1424. Providence, RI: American Mathematical Society.
- Zeh, H. (1973). Toward a quantum theory of observation. *Foundations of Physics* 3, 109–116.
- Zeh, H. D. (1993). There are no quantum jumps, nor are there particles! *Physics Letters A* 172, 189.
- Zurek, W. H. (1991). Decoherence and the transition from quantum to classical. *Physics Today* 43, 36–44. Revised version available online at <http://arxiv.org/abs/quant-ph/0306072>.
- Zurek, W. H. (1998). Decoherence, einselection, and the quantum origins of the classical: the rough guide. *Philosophical Transactions of the Royal Society of London A* 356, 1793–1820.