

Simulation of Condensed Gasses by Variation of their Coherent State Expansions

Rodney E. S. Polkinghorne



Presented on 10th May 2018 in fulfillment of
the requirements of the degree of Doctor of Philosophy

Abstract

The matter in a weakly interacting Bose-Einstein condensate is very nearly a classical wave. It can be accurately described by a mean field that follows a classical law of motion, with a small quantum field added that includes vacuum noise and other effects due to the atomic nature of matter.

The ground state of free bosons is a totally coherent field, but the ground state of interacting atoms is not totally coherent. If interacting atoms were prepared in a coherent state with the same density as their actual ground state, the density would remain static. However, the correlations in the density would oscillate, and this would give an observable difference between the two states.

The oscillating correlations can be computed by the phase space methods of quantum optics. In order to run efficient Monte-Carlo simulations, where the observables of experimental interest correspond to moments of the samples, it is almost always necessary to use the Wigner distribution, and to replace its law of motion by the classical law. This is a good approximation for systems that are close to classical wave dynamics.

Other systems can not be approximated classically in this way. These include processes of dynamical tunnelling, when the atomic field tunnels through amplitudes that a classical field could not have. Experiments are now proposed in which this will happen, and the decay of an atomic field by tunnelling from a metastable state is of widespread interest. To analyse such processes will require methods that allow for interference in superpositions of fields with different amplitudes. The obvious method, involving variational dynamics of coherent states, is analysed and tested. The numerical instability that has limited its use previously is analysed, and a way to resolve it proposed. But it is shown that this method, in its current form, can not simulate the dynamics of fields with high-dimensional classical phase spaces. The fair sampling assumption of the Monte-Carlo methods is essential to represent states of such fields.

Acknowledgements

First, I must thank my supervisors. Peter Drummond snuck me in to the atomic physics group at Swinburne and gave me a second chance at doing research. He has contributed useful ideas at some critical points in this project. Hui Hu has been a very helpful guide to the condensed matter literature, and the efficient research method of guessing who might have found the answer before I even thought of the question.

The other members of the theory group, especially Tim Vaughan, Bogdan Opanchuk, Laura Rosales-Zarate, Run Yan Teh, King Lun Ng, Simon Kiesewetter, Bryan Dalton, Qiong Yi He, Margaret Reid and Xia-Ji Liu, have been a very helpful source of idea and feedback, and very welcome company in our part of the scientific frontier.

I have received material support from a Swinburne University of Technology scholarship, and from Centrelink during the research assistantship that preceded my candidature. ACQAO and Ken Baldwin paid to send me to ICAP and the associated summer school.

Peter Hannaford and Russell McLean, the professors who run CQOS, have given me some very valuable advice and opened doors that made everything much easier. And I must thank Tatiana Tchernova, the department secretary who runs Peter and Russell. She is the life preserver to whom we all cling in the unplumbable sea of university bureaucracy.

My partner Patrick Ong, and my parents, have generously supported and put up me through all of this.

Figures 16 and 17 are reproduced with the permission of The American Association for the Advancement of Science.

Declaration

The work presented in this thesis has been carried out in the Centre for Quantum and Optical Science (CQOS) at Swinburne University of Technology in Melbourne between May 2011 and October 2016. The thesis contains no material that has been accepted for the award of any other degree or diploma. To the best of my knowledge, the thesis contains no material previously published or written by another author, except where due reference is made in the text of the thesis.

A handwritten signature in blue ink that reads "R Polkinghorne".

Rodney Polkinghorne
Hawthorn, Victoria, Australia
10th May 2018

Table of Contents

	List of Figures	7
1.	Introduction	9
1.	Atom optics	9
2.	Discretisation and Simulation	10
2.	Background	13
3.	Dispersion units for wave functions	13
4.	Coherent States	15
5.	Cold bosons	19
6.	Quantum fields	22
7.	Motion of the order parameter	24
8.	Geometry of the order parameter	26
9.	Graphs of wave functions	28
10.	Condition and stability	29
11.	Stiffness in quantum dynamics	31
12.	Frames	34
13.	Fourier series as frames	38
3.	Phase space Monte-Carlo methods	43
14.	Quartic oscillator	44
15.	Truncated Wigner dynamics	47
16.	Cahill frames for operators	53
17.	Phase space dynamics	55
18.	Positive P representations	57
19.	Monte Carlo dynamics	60
20.	Chemistry methods	61
21.	Summary	62
4.	Split condensates	65
22.	Split trap	66
23.	Gas microscope	70
24.	Bisection as beam splitting	71
25.	Models of the Split Trap	73
26.	Bogoliubov ground states	76

27.	Number fluctuation observables as Wigner distributions	85
28.	Oscillating correlations from a coherent state	87
29.	Results of Wigner simulations	88
30.	Correlations between numbers oscillating in a trap	98
31.	Wigner moments for interference observables	100
32.	Hartree modes	102
33.	One-dimensional field dynamics	105
34.	Sampling for Hartree modes	108
35.	Sound wave dynamics	110
36.	Field theory	113
5.	Expansions over Coherent States	116
37.	Multivariate derivatives	118
38.	Notation for kets and vectors	120
39.	How the variational method is supposed to work	121
40.	Expansions over frames	124
41.	Extrapolation in phase space	127
42.	Density of coherent amplitudes	132
43.	Discrete variational frame	134
44.	Regularisation	135
45.	Ill-posedness	137
46.	Tikhonov conditioning	140
47.	Effects of Tikhonov conditioning	143
48.	Conditioning with constraints	145
49.	Other approximants	147
6.	Conclusion	149
	References	151

List of Figures

1.	Length x and energy K^2 scales of a harmonic oscillator. Details in text.	17
2.	Multiplicity of condensed and non-condensed states.	20
3.	The colour scale used to represent complex phase.	28
4.	Examples of phase space coefficients.	29
5.	Eigenmodes with large frequencies become unstable when time is discretised.	34
6.	Example of frame bounds for a basis.	36
7.	Example of a redundant frame.	37
8.	Oversampled sinc wave packets in reciprocal space.	41
9.	The effect of the frequency of the wave packets.	42
10.	Phase space components of a coherent state evolving in a quartic potential.	46
11.	Area in phase space on which the components are large.	47
12.	Wigner distribution for a coherent state evolving in a quartic potential.	50
13.	Truncated Wigner dynamics for a coherent state evolving in a quartic potential.	52
14.	Modulus and phase of phase space components of cat states.	60
15.	Length, time, and energy scales for the Vienna split trap.	67
16.	Interference fringes from two coherent atomic clouds.	68
17.	Interference fringes from two incoherent atomic clouds.	69
18.	Delayed interference fringes from two coherent atomic clouds.	69
19.	Orbitals and corrections for repulsion.	74
20.	Amplitudes and energies of Bogoliubov modes.	83
21.	Effect of gas size on the coherence oscillations.	88
22.	Comparison of initial states, $\gamma = 0.01$, length 200, $\xi = 7.1$, $\xi^2 = 50$.	93
23.	Comparison of initial states, $\gamma = 0.03$, length 200, $\xi = 4.1$, $\xi^2 = 14$.	94
24.	Comparison of initial states, $\gamma = 0.1$, length 200, $\xi = 2.2$ and $\xi^2 = 5$.	95
25.	Oscillating coherence with $\gamma = 0.03$, length 20, $\xi = 4.1$, $\xi^2 = 14$.	96
26.	Oscillating coherence with $\gamma = 0.03$, length 10, $\xi = 4.1$, $\xi^2 = 14$.	97
27.	Momentum density by Bogoliubov and ABACUS.	98
28.	Oscillations in numbers of atoms on either side of a splitting trap.	99
29.	The Dirac-Frenkel method in parameter space and Hilbert space.	124
30.	Expansion of a coherent state over a set of 9 coherent states.	127
31.	Singular kets of the expansion operator over a disk in phase space.	131
32.	Singular values of the expansion operator over a disk in phase space.	131

33.	A square grid of coherent states covers phase space.	132
34.	A rectangular grid of coherent states does not overlap all of phase space.	133
35.	Residuals and coefficient sizes for variational frame expansions.	135
36.	Schematic diagrams of signal and noise in an ill-posed problem.	138
37.	Regularisation L-curve for a normal ill-posed problem.	141
38.	Cat state sampling problem.	142
39.	L-curve for cat state sampling problem.	143
40.	Picard plot for cat state sampling problem.	143
41.	Effect of Tikhonov regularisation on singular values.	145

This thesis is about computing the dynamics of the atomic fields of cold bosons. It approaches the subject from several different perspectives. Part of the work, described in Chapter 4, is fairly concrete, and uses a simple and well-understood truncated Wigner method to compute some spatial coherence properties of degenerate one-dimensional Bose gasses. Calculations of this type are ultimately limited by the inherent limitations of the method, as is discussed when it is derived in Chapter 3. One approach to working around these limitations is developed in Chapter 5. Although this approach has been used with some success in quantum chemistry, it turns out to have some basic restrictions that prevent it from being applied to quantum fields, at least in the way it was originally conceived.

1. Atom optics

In classical physics, matter is made of particles. This is one limit of quantum mechanics; it typically occurs at low densities and high temperatures, where the wave functions of different particles do not overlap. For matter in a gaseous phase, this limit is a classical ideal gas. Atom optics is the opposite limit. It occurs in systems of bosons at high densities and low temperatures, whose wave functions are nearly identical. In the atom optical limit, the atomic nature of matter is of little consequence; for many purposes, a cold and dense Bose gas can be treated as a classical wave in a field called the order parameter.

The technology to prepare a Bose gas in the atom optics regime was developed in the 1980s and 1990s, and is now quite routine [1]. The prerequisites were vacuum technology and the laser: laser cooling and trapping allow a gas to be prepared at very low temperature, and in a sufficiently rarefied vacuum it can be isolated well enough that it remains that way. The final invention was evaporative cooling, which created clouds of gas so cold that the wave functions of their atoms overlap. These are atomic Bose-Einstein condensates.

The atom optical limit of classical waves is very different from the limit of classical particles, but it retains the virtue of being simple to analyse. Quantum mechanics is based on the superposition principle; in a quantum system, the number of possible superposed configurations increases combinatorially with the number of moving parts. This complexity results from all the different ways that the state of each part can be correlated, or entangled, with the states of the other parts.

In the particle limit, these correlations do not occur, because the system is in a Hartree

state where the state of each particle is independent and well defined. This makes the system very much easier to analyse, and one of the most commonly used approximations in quantum mechanics is to assume that the state of a system has a Hartree form. Something very similar occurs in the wave limit, when the system is treated in second quantisation. When a coherent field is expanded over orbitals, its state factorises, and the state of each orbital can be described by a single complex amplitude. This form is much simpler than an arbitrary state of a field, where particles can occupy the orbitals with arbitrary correlations in their numbers.

This simplicity makes atomic Bose-Einstein condensates a very useful system for exploring the dynamics of quantum fields. Most of the time, it is possible to understand and predict the dynamics of a condensate at least qualitatively in terms of the classical field theory of its order parameter. It often suffices to think in even simpler terms, where the condensate is a classical fluid with an extra pressure term to account for the zero point energy. Also, the environment of a condensate is an atom trap, and the lasers, magnetic fields and microwave beams that form the trap can be controlled very precisely. This allows many condensed matter phenomena to be simulated by configuring an atom trap analogous to the material system in which they occur. By controlling the trap, and seeing how the phenomenon changes, it is sometimes possible to resolve questions for which the evidence provided by the material systems that exist in nature is ambiguous.

While many cold boson systems can be adequately analysed by treating the atomic field as totally coherent, others can not. These are often experiments that involve fluctuations in atom numbers or limited coherence times and lengths. The atomic nature of matter does not appear in the mean field theory of the order parameter, but it does appear in the fluctuations around the mean field. The discrete atoms cause shot noise when the density of a condensate is measured. Also, condensates confined to a line instead of a three-dimensional cloud have a finite coherence length, which requires a quantum treatment to understand.

These systems are not totally coherent, so they can not be treated simply with a mean field theory. However, the fact that they are nearly coherent can be exploited in order to analyse them more efficiently. This is the regime of quantum optics, where the uncertainty principle prevents the coherent amplitude of any system from being precisely defined, but the limited spread of the coherent amplitudes in phase space allows the states of the systems to be represented efficiently. This thesis uses quantum optical phase space methods to investigate some atom optical systems with limited coherence.

2. Discretisation and Simulation

In order to simulate a quantum system on a classical computer, the continuous wave function must be reduced to a finite set of numbers. At first glance, this seems like a major problem. A classical particle is described by two numbers, coordinates in phase space, but it takes a wave function, a continuously infinite set of numbers, to describe a quantum particle. The leap in the range of possibilities for the particle, and therefore in the information

required to describe it, looks overwhelming.

This impression is deceptive. Particles are generally found in a bounded region of space, and with a finite energy, so that they are constrained to a bounded region of phase space. For particles that require a quantum mechanical analysis, this region tends to be small, within a few orders of magnitude of Planck's constant. And the set of wave functions with restricted position and energy has Hilbert-space dimension roughly equal to the area of phase space it occupies, in units of Planck's constant [2].

Larger bounds are required to describe a particle more precisely, but in practice the required area is not too great. The lowest eigenstates of a harmonic oscillator can be calculated to ten significant figures on a phase space of area $30\hbar$, on which a wave function can be specified by a set of 30 complex coefficients [3]. This will be discussed in depth in the final chapter of the thesis, but the number of independent states available to a single quantum particle is never very large.

On the other hand, the superposition principle has interesting implications for systems of many particles. Since one particle can be in multiple places at a time, a set of them can be in combinatorically many configurations at once, and a vast number of probability amplitudes are required to specify a state. Traditionally, this has been considered a huge challenge. Great ingenuity has gone into inventing approximations, such as Dyson series, phase-space expansions and density-matrix renormalisation group (DMRG) algorithms, that approximate the states accessible to systems of interest while reducing the problem to a manageable number of parameters.

More recently, the combinatorial growth of configurations has been seen as an opportunity, sending legions of engineers on a quest for the quantum computer that is about to enter its third decade. From a scientific point of view, the phenomenon of entanglement has shown that, when several particles are involved, the consequences of the superposition principle are even stranger than they appear, and have made possible experiments that exclude the possibility of a predetermined reality to which quantum superpositions are an approximation.

Cold bosons are very simple in principle. A model of identical point particles with contact interactions almost always suffices. However, calculating properties of many-particle quantum systems is notoriously difficult in practice, and various methods have been applied to cold bosons. Green's function methods have a long history in condensed matter problems, being applied to calculate properties of systems such as superfluids and superconductors [4]. In lower dimensions, there are exact solutions for Bose gasses at zero temperature [5] and at finite temperature [6]. These have allowed some properties of cold Bose gasses to be calculated in closed form [7].

Another model for a one-dimensional Bose gas comes by identifying it with a Tomonaga-Luttinger liquid [8]. This is a universal model of a one-dimensional field in the limit of low energy, which is valid for a Bose or Fermi field, at weak or strong repulsion.

Computer simulations of the atomic field have been done in a number of ways. The simplest method, used from the earliest days of Bose-Einstein condensates, is to integrate

the Gross-Pitaevskii equation, a classical mean field theory [9, 10]. Methods based on the density-matrix renormalization group have been applied to many one-dimensional quantum systems, including cold atoms [11]. The other major approach has been the phase space methods that were invented to describe lasers, which have been adapted to treat atoms instead of photons [12].

A separate line of development is the Gaussian wave packet methods of quantum chemistry. These are related to the phase space methods of quantum optics, in a way that is examined in Chapter 5. The method was first proposed by Heller [13, 14] in the 1970s. It was developed over the decades since [15, 16, 17], and most recently in [18]. This will be developed further in Chapter 5.

The first purpose of this chapter is to explain some ideas that might not be encountered in a typical physics degree, but that later chapters take for granted. Section 4 derives some basic properties of the coherent states on which phase space simulation methods are based. Section 5 reviews what is known about the coherence of the atomic field along a gas of cold bosons that is confined to a line. Sections 10 and 11 discuss some properties of differential equations that hinder their numerical solution, and Sections 12 and 13 introduce a general way to analyse expansions of vectors over non-orthogonal sets such as the coherent states.

The second purpose is to introduce some notation and conventions that are used in the thesis. The analysis of variational methods in Chapter 5 requires the derivatives of some higher dimensional functions to be manipulated in moderately complicated ways, and these manipulations are more clearly expressed in the notation that mathematicians use for such derivatives. The numerical simulations in Chapter 4 use a scaled system of units, in which all quantities related to a condensate reduce to lengths. This restricts the magnitudes of the quantities to a range that computers can handle. It is also quite helpful when reasoning physically about cold atoms, and about how the geometry of a trap governs their mechanics.

3. Dispersion units for wave functions

Describing an atom trap in detail requires electromagnetism, and units for charges and fields. But that is outside the scope of this thesis: for our purposes, mechanics suffices. The trap interests us only because it applies a mechanical potential to atoms, and we will describe that potential and the atomic motion in terms of quantities derived from length, time and mass. These quantities are not convenient to express in the International System of Units, nor any system that stems from a human scale, because the atoms are too small and move too fast. We will derive our units from their scale instead of from ours [19].

The most useful scale that an atom defines is its mass. Provided that only one isotope is present, this mass is the only one that atom optics involves. Since mass enters into quantum mechanics through the dispersion relation for de Broglie waves, this characteristic mass allows wavelengths to be identified with periods. In the process, every occurrence of the atomic mass can be absorbed into the units.

The resulting system of dispersion units has only one dimension, whose unit can be either a length or a time. In cold atoms, it is usually convenient to use a length. Other experimental systems, such as magnetic resonance, involve controlling how parameters

vary over time in sophisticated ways. This is rarely done in an atom trap; when the trap geometry changes, it happens either slowly enough that the atoms follow it adiabatically, or fast enough to be instantaneous compared to their motion. The relevant times are now and never, which need no units. On the other hand, the geometry of the trap sets many physically meaningful lengths.

Condensates will be described formally in Section 8, and they will turn out to have two natural length scales. The order parameter that defines the condensate can often be treated as a classical field, whose motion is determined by a scale called the healing length. When quantum fluctuations are added to this classical description, those fluctuations are due to shot noise, and the scale on which they become significant is the separation between atoms. The ratio of these lengths is called the Lieb-Liniger parameter, and it measures how far the condensate is from being a classical field. When possible, it helps to select one of these lengths as a natural unit. However, this thesis considers highly anisotropic condensates, whose axes have very different scales. This makes it hard to define the atomic separation and the healing length consistently, and it is often convenient to use the micrometre instead.

To derive these dispersion units, we start with Schrödinger's equation for a wave function ψ in conventional mechanical units,

$$i\hbar\partial_t\psi = H_{\text{SI}}\psi = \left(-\frac{\hbar^2}{2m}\nabla^2 + V\right)\psi. \quad (1)$$

Each side has the dimension $\text{kg}\cdot\text{m}^{1/2}\cdot\text{s}^{-2}$, so this depends on all three mechanical units. We would like to identify lengths with times as we do in relativity, but there is a complication: the speed of light is fixed, but the speed of an atom varies. De Broglie waves satisfy a quadratic dispersion relation $\omega = (\hbar/2m)k^2$, so a time t must be identified with an area $s = (\hbar/m)t$ instead of with a length. The group velocity is $d\omega/dk = (\hbar/m)k$, and the trajectory of an atom moving at that speed satisfies $(\hbar/m)k = x/t$, whence $k = x/s$. Dispersion units identify the speed of an atom x/s with its de Broglie wave number k . Frequency is identified with reciprocal area; we will use the symbol χ in place of ω to make it clear that the units have been scaled.

The relation between functions of conventional time and the equivalent functions of scaled time is $f_{\text{SI}}(t) = f(\hbar t/m)$, whence $\partial_t f_{\text{SI}} = \hbar/m \partial_t f$. Some equations, such as Equation 1, are explicitly stated to be in conventional units. This means that all functions of time that occur in them are the SI versions. Some constant quantities involve units of time, Hamiltonians for instance, and these are given an explicit SI subscript.

When the quantity s is substituted for t , so that the wave function is $\psi(s, x, y, z)$, Equation 1 reduces to

$$2i\partial_s\psi = 2H\psi = (-\nabla^2 + K^2)\psi. \quad (2)$$

The quantity V in Equation 1 is an external potential, usually that of the atom trap. In Equation 2, $K^2 = 2mV/\hbar^2$ is the square of the wave number of an atom with that much kinetic energy. The conversion factor between V and K^2 is the last place that the atomic mass m occurs. This form makes it apparent that the natural units for Schrödinger's equation must

give ∂_s the same dimensions as ∇^2 , and hence they must identify time with area instead of with length. The units of ψ have not changed, and wave functions that were normalised in SI units remain normalised in the new dispersion units. This is another advantage of retaining the dimension of length instead of time.

Next, consider eigenstates. Suppose that a wave function ϕ satisfies Equation 1, with

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + V\right)\phi = E\phi \quad \text{and} \quad \partial_t\phi = -i\frac{E}{\hbar}\phi.$$

The corresponding solution to Equation 2 satisfies the scaled relations

$$(-\nabla^2 + K^2)\phi = k^2\phi \quad \text{and} \quad \partial_s\phi = -i\frac{k^2}{2}\phi.$$

Here, $k^2 = 2mE/\hbar^2$, with the same conversion that relates K^2 to V . Note that wave number relates to scaled frequency as $\chi = k^2/2$, retaining the factor of 2 from the de Broglie dispersion relation.

4. Coherent States

It is instructive to apply the scaled units of Section 3 to a simple harmonic oscillator. The results are used below to analyse harmonic atom traps. More importantly, bosonic fields are identical to systems of harmonic oscillators [4]. The simplest description for the field of a condensed gas of bosons is one of the coherent states, the set of kets that comprises the ground state of a harmonic oscillator and the states that result from translating and accelerating it away from the origin of phase space.

Atom optics is concerned with coherent atomic fields. These are a relatively new phenomenon, but coherent states of light have been studied in great detail since the invention of the laser [20]. The dynamics of light waves, and of condensed bosonic atoms, are equivalent to a set of harmonic oscillators. When the state of the field is fully coherent, the amplitudes of these oscillators are defined to the uncertainty limit. Such states are known as coherent states.

Coherent states were discovered very early in the history of quantum mechanics [21]. The idea that measurements are essentially uncertain prompted the first quantum theorists to consider physical states that cause measured values to be normally distributed, and to look at Gaussian wave packets. These turn out to be the quantum states whose motion in a quadratic potential best approximates simple harmonic oscillation. A quadratic potential sets a natural width for wave packets, determined by equipartition of energy and the uncertainty principle, and packets maintain this width as they oscillate. A wave packet with this width is called a coherent state.

Coherent states were first considered as wave packets in first quantisation, but they have been most useful in second quantisation, where they are states of atomic fields. The invention of the laser in 1963 called for a quantum theory of intense radiation to describe it. The first problem was the size of the quantum description of a laser beam, up to 10^{18}

photons distributed across thousands of electromagnetic modes. Such states are elements of a Hilbert space with very high dimension, and it is not practical to record them, either by hand or in computer memory. Further progress required an efficient way to approximate them.

This was found by Glauber [22] and Sudarshan [23], who showed that any quantum state can be expanded as a superposition of coherent states. Laser light is very coherent, so these expansions often represent it efficiently, converging with a reasonable number of states in the superposition. Moyal found an efficient way to calculate the dynamics of these superpositions [24], and a family of phase space Monte-Carlo methods developed from his idea. When very coherent matter became available in the lab, in the form of dilute gas Bose-Einstein condensates, these methods were adapted to calculate its dynamics. They have since been extended to other forms of condensed matter, solving many problems, but they are limited by intractable convergence issues.

As well as representing a fully coherent field, coherent states are used in computer simulations, through the method of phase space expansions. We will use these methods extensively in this thesis. Phase space is a plane whose axes are a pair of canonical coordinates. When the coordinates are those of a harmonic oscillator, it makes sense to think of this as a complex plane. The complex numbers points are phasors, or complex amplitudes. By convention, the real axis is the position coordinate, and the imaginary axis is the momentum coordinate.

In conventional mechanical units, let the resonance frequency of an oscillator be ω_0 , and the Hamiltonian be

$$H_{\text{SI}} = \frac{p^2}{2m} + \frac{1}{2}m\omega_0^2x^2. \quad (3)$$

In dispersion units, this scales to

$$2H = l_0^{-2} \left(l_0^2 k^2 + \frac{x^2}{l_0^2} \right) = 2l_0^{-2} (a^\dagger a + \frac{1}{2}),$$

where $l_0 = (\hbar/m\omega_0)^{\frac{1}{2}}$ is the natural length scale of the oscillator, and the annihilation operator is

$$a = \frac{1}{\sqrt{2}} \left(\frac{x}{l_0} + \frac{il_0 p}{\hbar} \right). \quad (4)$$

The vacuum energy $\frac{1}{2}\hbar\omega_0$ scales to a wave number l_0^{-2} . Figure 1 shows the atomic densities of the ground state ψ_0 and the first excited state ψ_1 , their energies, and the position uncertainty of the ground state. The momentum uncertainty scales to a wave number $\Delta k = (\sqrt{2}l_0)^{-1}$, and the scaled uncertainty principle is the Fourier transform law $\Delta x \Delta k \geq \frac{1}{2}$.

The coherent state $|\alpha\rangle$, an eigenstate of the annihilation operator defined in Equation 4, satisfies $a|\alpha\rangle = \alpha|\alpha\rangle$. The operator a is not hermitian, so α is a complex number. The unit area defined by x and k coordinates differs by a factor of 2 from the unit area of the complex plane: the relation is $dx dk = 2d^2\alpha$.




Figure 1. Length x and energy K^2 scales of a harmonic oscillator. Details in text.

The coherent state $|0\rangle$ satisfies $a|0\rangle = 0$, so it is the ground state of the harmonic oscillator. The other coherent states are displacements of this to nonzero amplitudes, and their wave functions are Gaussian wave packets. Let $\alpha = (x_0/l_0 + il_0k_0)/\sqrt{2}$. The wave function for a coherent state is

$$\langle x|\alpha\rangle = \frac{1}{\pi^{1/4}l_0^{1/2}} \exp\left(-\frac{(x-x_0)^2}{l_0^2} + ik_0x\right). \quad (5)$$

Coherent states are not orthogonal, but their brackets follow the law

$$\langle\beta|\alpha\rangle = e^{-\frac{1}{2}|\beta|^2 - \frac{1}{2}|\alpha|^2 + \beta^*\alpha} = e^{-\frac{1}{2}|\beta-\alpha|^2 + i\text{Im}\alpha^*\beta}, \quad (6)$$

whence

$$|\langle\beta|\alpha\rangle|^2 = e^{-|\beta-\alpha|^2}. \quad (7)$$

When the kets are considered as vectors in Hilbert space, the angle between them satisfies

$$\cos\theta = \frac{|\langle\beta|\alpha\rangle|}{\|\alpha\|\|\beta\|} = e^{-\frac{1}{2}|\beta-\alpha|^2},$$

and comparing Taylor series gives

$$\theta = |\beta - \alpha| + O(|\beta - \alpha|^2).$$

The states are directions in Hilbert space, and, when the angles between these directions are small, the phase space plane maps them without distortion.

In classical mechanics, phase space is the plane of canonical coordinate pairs (q, p) . This can equally well be considered as a complex plane, where canonical coordinates are

represented by complex numbers α . The unit of area is an action, so the scale of classical phase space is arbitrary.

Quantum phase space is also a complex plane, but its points are coherent states instead of classical states of motion. This means that it has quite a different geometry. Locally, it maps onto a plane in Hilbert space without distortion, and has the same geometry as classical phase space, except that angles in Hilbert space do not correspond to angles in the plane, but rather to distances, and solid angles correspond to areas. Quantum phase space has a natural unit of area, derived from the uncertainty principle, and a circle with this area gives a precise definition of local. Over distances large compared to the uncertainty scale, the coherent states are very nearly orthogonal.

The projectors onto the coherent states satisfy a relation

$$\int |\alpha\rangle\langle\alpha| d^2\alpha = \pi. \quad (8)$$

A set that satisfies a relation of this form is called a tight frame. If the right hand side were 1, the coherent states would be an orthonormal basis of ket space. Frames are a generalisation of orthonormal bases, which we will discuss further in Section 12. Any state can be expanded over coherent states, in the form

$$|\psi\rangle = \pi^{-1} \int |\alpha\rangle\langle\alpha|\psi\rangle d^2\alpha. \quad (9)$$

However, this expansion is not unique: a coherent state $|\beta\rangle$ has the trivial expansion $|\beta\rangle = |\alpha\rangle$, but the expansion given by Equations 6 and 9 is a Gaussian function of α .

This thesis will often mention the derivative of a coherent state with respect to its coherent amplitude. To facilitate such discussion, we will define a function $|o\rangle$ from the complex numbers into ket space, where $|o(\alpha)\rangle = |\alpha\rangle$. This is not an analytic function, in the sense that there is no function $|Do\rangle$ such that $|o(\alpha+h)\rangle = |o(\alpha)\rangle + |Do(\alpha)\rangle \cdot h + o(h^2)$. However, the function of real variables $(x, y) \rightarrow |o(x+iy)\rangle$ is differentiable. Equivalently, there are Wirtinger derivatives $|Do\rangle$ and $|D^*o\rangle$ such that $|o(\alpha+h)\rangle = |o(\alpha)\rangle + |Do(\alpha)\rangle h + |D^*o(\alpha)\rangle h^* + o(h^2)$. These Wirtinger derivatives are often easier to manipulate.

The theory in this thesis will be done almost purely in terms of coherent states, but some of the numerical experiments will expand these over harmonic oscillator number states for simplicity of implementation. The expansion is based on the formulae for the displacement operator

$$D(\alpha) = e^{\alpha a^\dagger + \alpha^* a} = e^{-\frac{1}{2}|\alpha|^2} e^{\alpha a^\dagger} e^{\alpha^* a} = e^{\frac{1}{2}|\alpha|^2} e^{\alpha^* a} e^{\alpha a^\dagger}.$$

Applying the first equality to the vacuum state gives

$$|\alpha\rangle = D(\alpha)|0\rangle = e^{-\frac{1}{2}|\alpha|^2} e^{\alpha a^\dagger} |0\rangle = e^{-\frac{1}{2}|\alpha|^2} \sum_n \frac{\alpha^n}{\sqrt{n!}} |n\rangle. \quad (10)$$

The state of a system that is not a harmonic oscillator can still be expanded over coherent states, using a suitable l_0 to provide length and momentum scales. This can be an efficient representation when the Hamiltonian is almost harmonic. Further properties of the coherent states will be stated as required. Glauber summarised these in a classic paper [22], which has been paraphrased in every quantum optics textbook written since.

5. Cold bosons

This thesis is largely about techniques for numerical simulations; the physical system that inspires those simulations is Bose-Einstein condensation in dilute gases [1]. The reader is assumed to have heard of this, the pin-up phase of matter of the 1990s, and to have some idea why it is important. This section reviews some of the properties that are important for this thesis.

The condensates of interest for this thesis, that occur as a homogeneous and isolated system, are an entirely artificial phase of matter. They comprise a thousand to a few million atoms in a rarefied vacuum chamber at a temperature below a nanokelvin. These are held in place and cooled by various ingenious schemes involving lasers, microwaves and magnetic fields; as mentioned in Section 3, this thesis will take those mechanisms for granted, supposing that they happen quickly compared to the motion of each atom, which is governed by an effective adiabatic potential.

There is another type of quantum gas, the degenerate Fermi gas, in which the particles are fermions instead of bosons. This thesis will not consider those directly. Some of the simulation techniques are potentially applicable to them, but there are significant technical difficulties in applying them [25].

The major difference that quantum mechanics makes to the properties of a gas is that quantum particles are identical, while classical mechanics assumes that all particles can in principle be distinguished [26]. The point at which the classical and quantum descriptions diverge is their combinatorics. When a set of identical bosons is distributed among a set of orbitals, the number of possible arrangements is less than it would be if the particles were distinguishable, as illustrated in Figure 2. Also, and this is the critical point, the missing arrangements are ones in which the bosons are spread over many orbitals. When the particles are fully spread out, each of them having its own orbital, there are $N!$ arrangements of the distinguishable particles, and only one of the identical bosons.

The consequence of this is that bosons stick together in their ground orbital at much higher temperatures than distinguishable particles do. When a set of bosons is in thermal contact with a bath, the sharing of energy between the bosons and the bath is determined by combinatorics. Giving extra energy to the bosons allows them to access more arrangements. The temperature is defined by $dU = \tau ds$ —we will use the units where temperature is an energy $\tau = k_B T$ and entropy is the dimensionless logarithm of multiplicity—so temperature is inversely related to the rate at which these extra arrangements become available. For a given arrangement of particles, this rate is lower for bosons than it is for distinguishable particles, therefore the bosons arrange themselves that way at a higher temperature.

Quantitatively, classical particles in contact with a bath follow a Boltzmann probability distribution, where the probability of them being arranged with energy ϵ is proportional to

$$P(\epsilon) = e^{-\epsilon/\tau},$$

but identical bosons follow

$$P(\epsilon) = \frac{1}{\lambda^{-1}e^{\epsilon/\tau} - 1},$$

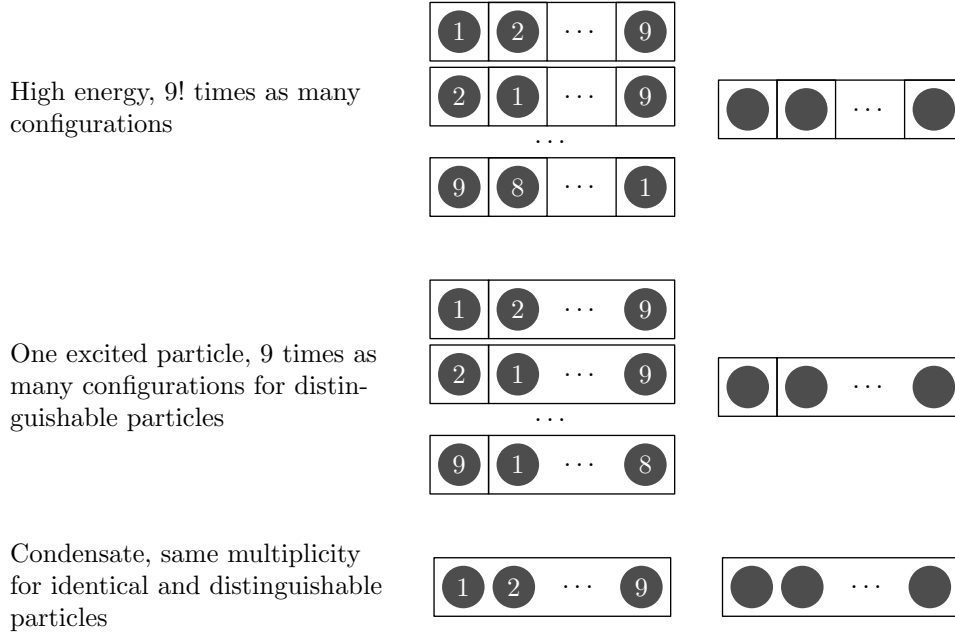


Figure 2. Multiplicity of condensed and non-condensed states.

the Bose-Einstein distribution, with $\lambda = \exp(\mu/\tau)$ being the ratio of chemical potential to temperature. In order for this probability to be positive, the chemical potential μ must be less than the energy of the lowest orbital. We will consider an ideal gas, whose only energy is kinetic, and whose ground orbital has zero energy. In this case, for condensation to occur and $P(0)$ to be large, the number λ must be very slightly less than one.

Now, following [27], consider a gas of free bosons, in a large box of volume L^3 . Wave numbers are quantised with $k = 2\pi n/L$, where n is a vector of three integers. The relation between index and energy is $\epsilon = 4\pi^2 n^2/L^2$, and the number of orbitals of energy ϵ or less is

$$N = \frac{4\pi}{3} n^3 = \frac{2}{3} \frac{L^3}{(2\pi)^2} \epsilon^{3/2},$$

so the density of orbitals is

$$\frac{dN}{d\epsilon} = \frac{L^3}{4\pi^2} \epsilon^{\frac{1}{2}}.$$

The density of the gas is

$$n = \frac{1}{L^3} \int_0^\infty P(\epsilon) dN = \frac{1}{4\pi^2} \int_0^\infty \frac{\epsilon^{\frac{1}{2}}}{\lambda^{-1} e^{\epsilon/\tau} - 1} d\epsilon. \quad (11)$$

This actually counts the bosons in the excited orbitals, because the probability for the ground orbital, with $\epsilon = 0$, is multiplied by a zero density; when the particles are condensed, and a macroscopic number of particles are in that orbital, the integral in Equation 11 does not approximate the original sum over orbitals. However, the approximation is still valid for the excited states.

If we substitute $x = \epsilon/\tau$, and assume that λ is too close to 1 to affect the integral, the result is

$$n = \frac{1}{4\pi^2} \int_0^\infty \frac{\epsilon^{\frac{1}{2}}}{e^{\epsilon/\tau} - 1} = \frac{\tau^{3/2}}{4\pi^2} \int_0^\infty \frac{x^{\frac{1}{2}}}{e^x - 1} dx. \quad (12)$$

The integral evaluates to $\frac{1}{2}\sqrt{\pi}\zeta(3/2)$ where ζ is the Riemann zeta function. The key point is that it is finite, and, at a given temperature, only a limited number of bosons are in the excited states, no matter how nearly the chemical potential rises towards zero. This sets a critical relationship between density and temperature,

$$n = 0.05864 \times \tau^{3/2}.$$

Below this temperature, or above this density, a macroscopic fraction of the gas condenses in the ground orbital. This is the familiar idea, stated in dispersion units, that the critical atomic separation is around the de Broglie wavelength corresponding to the temperature.

To compare this to the case of distinguishable particles, consider a harmonic trap with characteristic length l_0 . The degeneracy temperature at which distinguishable particles will condense in the ground orbital is its energy l_0^{-2} . If there are N identical bosons in the trap, their density is around Nl_0^{-3} , so the critical temperature for condensation is around $N^{2/3}l_0^{-2}$, which can be much higher than the degeneracy temperature.

Cooling through the critical temperature is a phase transition. At higher temperatures, the bosons are a classical gas: they have different momenta, allowing them to be distinguished from each other. They are not distinguishable at lower temperatures: the gas in the condensate is similar to a classical wave, and the fact that it is composed of atoms is difficult to observe.

The number of atoms that are condensed in the ground orbital is an order parameter for the phase transition. It tends to zero as the condensate heats and approaches the critical temperature. By convention, this number is expressed as a function of position ψ , proportional to the wave function for the ground orbital, and normalised so that $\int |\psi|^2$ is the number of atoms occupying it. If the condensate is disturbed, the atoms move collectively; there is still a macroscopically occupied orbital, but this will be an excited state and can vary with time. Its dynamics are discussed in Section 7.

A one-dimensional gas is different. The relation between ϵ and n^2 is the same, but now n is a single integer, whose value runs from $-N/2$ up to $N/2$. The number of particles with energy at most ϵ is given by $\epsilon = \pi^2 N^2/L^2$, so that $d\epsilon = 2\pi^2 N/L^2 dN = (2\pi/L)\sqrt{\epsilon} dN$, and $\rho(\epsilon) = (L/2\pi)\epsilon^{-\frac{1}{2}}$. The one dimensional equivalent of Equation 12 involves an integral

$$\int \frac{x^{-\frac{1}{2}}}{e^x - 1} dx$$

that diverges, because the integrand is close to $x^{-3/2}$ for small x . This means that, at any nonzero temperature, the excited orbitals can hold as many bosons as are present, and there is no phase transition.

At these low temperatures, gases are not ideal, and collisions between atoms can contribute significantly to their mechanics. The one-dimensional gas has two very different

limits at low temperature, depending on the size of this contribution. When collisions are important, the positions of the bosons correlate strongly in order to reduce the repulsion energy from their interactions: this is called the Tonks gas regime. The opposite limit, where the repulsion energy is dominated by the energy required to confine an atom within the atomic separation, is called a quasicondensate. In this limit, the wave functions of the atoms overlap, and the gas resembles a classical wave. The gas becomes uniform at low temperature, but the phase of the classical wave fluctuates on a scale [28]

$$\lambda_\tau = \frac{\hbar^2 n}{mk_B T} = 2n/\tau. \quad (13)$$

Chapter 4 is largely concerned with the spatial coherence of these quasicondensates.

6. Quantum fields

A gas of cold bosons is a system of many particles. Their momenta are limited by the lack of thermal energy, and their density is high enough that the area of phase space available per particle is less than \hbar , and the particles are physically indistinguishable. The simplest formalism to describe such particles is second quantisation [4]. Instead of describing the atoms as particles, this formalism describes the atomic field as a wave, in terms of a field operator $\hat{\Psi}(x)$. This satisfies a commutation relation $[\hat{\Psi}(x), \hat{\Psi}^\dagger(x')] = \delta(x - x')$, ensuring that the bosons occupying any orbital $\phi(x)$ can be counted by an observable of the form $n_\phi = a_\phi^\dagger a_\phi$, where $a_\phi = \int \hat{\Psi}(x)\phi(x) dx$.

The Hamiltonian that is used to describe a cold Bose gas, where the atoms attract or repel each other and are contained by a potential $V(x)$, is almost always of the form [1]

$$H_{\text{SI}} = \int -\frac{\hbar^2}{2m} \hat{\Psi}^\dagger \nabla^2 \hat{\Psi} + V \hat{\Psi}^\dagger \hat{\Psi} + \frac{2\pi\hbar^2 a}{m} \hat{\Psi}^{\dagger 2} \hat{\Psi}^2.$$

This Hamiltonian constrains the form of the atomic interactions. It describes atoms that interact pairwise, so that the interaction energy for atoms with positions x_i is a sum over terms $V_a(x_i - x_j)$. This is often adequate to describe an isolated gas, but the interactions that absorb atoms into the gas and scatter them out of it usually involve three atoms [1]. The range of the potential V_a needs to be shorter than the typical separation of the atoms, and the energy of the atoms needs to be small enough for their interaction to be described by an s-wave scattering length a . This is almost always the case when a condensate is formed in the laboratory. In dispersion units, the Hamiltonian becomes

$$2H = \int \hat{\Psi}^\dagger (-\nabla^2 + K^2) \hat{\Psi} + 4\pi a \hat{\Psi}^{\dagger 2} \hat{\Psi}^2. \quad (14)$$

When a quantum field is computed, it can be represented in either first or second quantisation. The field is described by a set of orbitals and a set of particles that occupy them; the efficiency of the two representations is governed by the ratio of particles to orbitals, which is roughly the same as the density of particles in phase space. A first quantised representation needs to store a wave function for each particle, which can be efficient when

the particles are spread out in phase space and the wave functions are distinct. A second quantised one records the occupations of a basis of orbitals that spans the states accessible to a particle. When the particles are clustered in phase space, this basis can have fewer elements than there are particles; this representation is especially efficient when the basis elements are represented by a few parameters instead of by sampled wave functions. A high density in phase space indicates that the system is more similar to a classical wave than a set of classical particles.

When a gas of bosons condenses, all of its particles occupy the same orbital. In this case, second quantisation is a very efficient representation: all that needs to be recorded is the order parameter that describes the occupied orbital. Computing the motion of the gas is quite simple in this representation; Section 7 derives a partial differential equation satisfied by the order parameter, called the Gross-Pitaevskii equation.

In second quantisation, the occupied orbital is equivalent to a harmonic oscillator whose quantum number is the number of bosons in the condensate. The state of this oscillator should strictly be a number state. However, calculations are much simpler if we pretend it is a coherent state. When we simulate quantum phenomena in the gas, atoms will scatter to orbitals other than the condensed one, and it is often convenient to expand the state over a fixed basis of orbitals, such as plane wave states, that does not include the orbital occupied by the condensate. It is simple and efficient to expand a coherent state in this way, because it remains a Hartree state under a unitary transformation of the orbitals, and expands to a product of coherent states in the oscillators that describe the second quantised orbitals. In contrast, a number state expands to a multinomial distribution of particles among orbitals, which requires a lot of bookkeeping to represent on a computer.

Coherent states have another advantage. Section 15 will introduce the truncated Wigner method that has been most successful at simulating wave-like quantum optical systems. This method starts by mapping the states of a quantum field to probability distributions over the phase space of the equivalent classical field, and proceeds to draw sample states from such a distribution to initialise a Monte-Carlo method. Under this mapping, the coherent states always map to probability distributions, but the number states do not: they map to functions that take negative values on some parts of phase space, and therefore can not be normalised to probability distributions. The coherent state combines these functions in a way that produces a positive distribution function.

Treating the condensate orbital as a coherent state makes it a superposition of number states with a range of occupation numbers, as in Equation 10. This is not physically correct. In principle, it is possible to prepare a condensate with an exactly known number of particles. If this were done, the results of some measurements would differ from the predictions for an atomic field in a multimode coherent state. For example, in a condensate with exactly 10^4 atoms, the variance of the particle number is 0, while it would be 100 for a coherent field. Some recent experiments at Harvard have demonstrated this [29].

In a coherent state, the number of particles in each orbital is not correlated with the numbers in the other orbitals. Physically, these numbers are entangled: if atoms occupy one

orbital, those atoms must be absent from the other orbitals. A coherent state predicts the distribution of particle numbers at each point very closely, but does not accurately model the correlations between many points, such as those that determine the variance in the total number of atoms.

However, the number of particles is conserved. When the coherent state is expanded over number states, each component propagates independently under time evolution. The dynamics of the components are entirely realistic, because each of them has a fixed number of atoms. Observables that are not sensitive to the total condensate size will have the same expectation value for each component, and this will be the expectation value for the coherent state. So for these observables—and, apart from a few experiments by Greiner [30] and Haroche [31], they are the only ones that can be measured—coherent states give the same predictions as number states.

7. Motion of the order parameter

Consider a condensed gas of bosons whose order parameter is a field $\psi(x)$, and an orthonormal basis of orbitals with wave functions $\phi_i(x)$. When the gas is represented in second quantisation, eigenstates of occupation number are identified with the stationary states of a simple harmonic oscillator, so that the occupation of each orbital is identified with the quantum state of a single particle; this allows the gas to be identified with a collection of distinguishable quantum particles labelled by the orbitals. As discussed in Section 6, the state of the gas is approximately coherent, an approximation that is convenient to work with and accurate enough to explain the types of experiments that are considered in this thesis.

The state of the condensed gas is a product of coherent states of the orbital occupations,

$$|\Psi\rangle = |\alpha_1\rangle \otimes |\alpha_2\rangle \otimes \cdots,$$

where $|\alpha_i\rangle$ is a coherent state that counts the occupation of the orbital ϕ_i . The amplitude α_i is the component of the order parameter in this orbital,

$$\alpha_i = \int \phi_i^* \psi.$$

It can be shown [20] that this state is defined consistently for any choice of basis ϕ_i . However, the result is especially simple when the condensed orbital is a element of the basis, say $\psi(x) = \sqrt{N}\phi_1(x)$, where N is the number of particles in the condensate. In this case $\alpha_1 = \sqrt{N}$, and all the other orbitals are in vacuum states, with $\alpha_i = 0$ for $i > 1$. On the other hand, for localised orbitals, where ϕ_i is a wave packet centred on the point x_i , the amplitude α_i will be close to the value $\psi(x_i)$ of the order parameter.

Now consider the trajectory of this coherent state under the Hamiltonian of Equation 14. The first term, $\int \hat{\Psi}^\dagger (-\nabla^2 + K^2) \hat{\Psi}$, will clearly preserve the coherent form: the basis ϕ_i can be chosen as eigenstates of the operator $-\nabla^2 + K^2$, so that the second quantised Hamiltonian is a set of independent harmonic oscillators, one for the occupation

state of each orbital. When the states of the oscillators are initially coherent, the trajectory generated by this simple harmonic Hamiltonian remains coherent. However, the term proportional to $(\hat{\Psi}^\dagger)^2 \hat{\Psi}^2$ becomes a quartic oscillator in second quantisation. The trajectory that this generates is described in Section 14; a coherent state does not remain coherent. The physical reason for this is that the kinetic and potential energies per particle are independent of the occupation of an orbital, but the repulsion energy per particle increases as more particles are absorbed into the gas.

When the gas is dilute, as is usually the case in laser trapping experiments, the repulsion energy is small enough that the dephasing due to repulsion can be ignored. In this case, a coherent state is a good approximation to the dynamics of the gas. Using the variational method that is described in Section 39, it can be shown [32] that the coherent amplitude of this state, which is the order parameter of the condensate, satisfies a differential equation

$$2i\partial_s\psi = (-\nabla^2 + K^2 + 8\pi a|\psi|^2)\psi. \quad (15)$$

This is the Gross-Pitaevskii equation. It is also the equation of motion for a classical field ψ , when the operator $\hat{\Psi}$ in the Hamiltonian of Equation 14 is replaced by a classical field amplitude ψ .

The procedure of assuming the state of a system is coherent, and applying the Dirac-Frenkel principle to derive an equation of motion for the coherent amplitude, was first used by Heller in 1976 [14]. Later, the method was extended to allow superpositions of several coherent states. Heller's work involved chemical reactions and the dynamics of small molecules with a few dozen degrees of freedom; Chapter 5 of this thesis will attempt to apply his method to quantum fields with many degrees of freedom, and will demonstrate why that is likely to work only in very special cases. Frenkel independently invented several methods that are very similar to the truncated Wigner method used to compute quantum dynamics, and we will mention him again in Section 15.

Although the order parameter is a classical field, the Gross-Pitaevskii Equation can still describe phenomena, such as atoms tunnelling through a potential barrier, that are usually considered to be quantum phenomena. In a first quantised picture, such tunnelling appears to be entirely quantum. The first-quantised classical limit is Hamiltonian mechanics, where each particle has a position and a momentum, and tunnelling is prohibited by conservation of energy. On the other hand, in a second quantised classical limit, the state of the particles is described by a collective wave function, whose classical field theory describes tunnelling in exactly the same way as Schrödinger's Equation does. This is not just splitting hairs. Photons are a type of particle that we are used to describing in second quantisation. No one talks about photons tunnelling, even though when evanescent waves and frustrated total internal reflection are considered from a first quantised point of view, that would be the obvious way to describe them. These are not considered to be quantum phenomena. The category of “classical” is an artefact of the way that theory developed historically, and it reflects the sequence in which phenomena were discovered and understood as much as it does any kind of internal consistency. Nature, of course, is consistently quantum.

8. Geometry of the order parameter

Much of this thesis is concerned with Bose-Einstein condensates. At the simplest level of approximation, a condensate comprises a large number of atoms that all have a common wave function. This is the classical wave limit, where we can forget that matter is made of atoms and treat the de Broglie wave as a classical field whose state is coherent. For now we will study the classical dynamics of this field, which completely explain many condensate phenomena. In Section 26 below, we will introduce Bogliubov modes, a more sophisticated approximation to the quantum state of a condensate and its dynamics that is sometimes more accurate.

The order parameter $\psi(s, x, y, z)$ satisfies Equation 15, the Gross-Pitaevskii equation. This would be a Schrödinger equation, except for the nonlinear term $8\pi a|\psi|^2\psi$ which describes repulsion between the atoms in terms of the s-wave scattering length a . The length ξ defined by $\xi^{-1/2} = 8\pi a|\psi|^2$ is called the healing length: apart from a numerical factor, its square is the ratio of the volume per particle to the scattering length. Equation 15 can be written

$$2i\partial_s\psi = (-\nabla^2 + K^2 + \xi^{-2})\psi.$$

but note that ξ differs across the cloud as the density of the gas does.

In the classical wave limit, the motion of the condensate is governed by Equation 15. In this limit, there is no way to tell that the atoms are discrete particles, so the atomic density $|\psi|^2$ and the repulsion energy per atom that is described by a do not have independent physical consequences: the motion of the gas is determined by their product, which determines the healing length. The other parameter, the external potential K^2 , describes the environment, while the healing length describes the intrinsic properties of the gas. In the Thomas-Fermi limit, where the repulsion energy dominates the zero-point energy of the trap, the energy ξ^{-2} is close to the chemical potential of the gas referred to the trap potential at that point of the condensate.

Equation 15 is satisfied by an untrapped gas in the form of a dark soliton $\psi(x, y, z) = \psi_0 \tanh(x/\Xi)$, where Ξ is the healing length far from the x -axis, satisfying $\Xi^{-2} = 8\pi a|\psi_0|^2$. This is the geometric meaning of the healing length. The density of the soliton falls to zero at the x -axis; this solution shows how the density of a condensate falls to zero at a hard edge. It is difficult to imagine a configuration where the condensate changes more sharply than it does at a hard edge, so the healing length usually sets the scale for the smallest features in the order parameter, and therefore in the classical description of the condensate.

Many properties of the condensate change qualitatively at the scale ξ , because the soliton occurs on the scale where kinetic energy balances repulsion. On shorter scales, the kinetic energy dominates and the mechanics resemble a free quantum particle; on longer scales, repulsion dominates and the motion is similar to fluid dynamics. The dispersion relation for sound waves in the condensate changes from quadratic de Broglie waves at long wavelengths to linear sound waves at short wavelengths, as discussed in Section 26.

This thesis will consider atoms that are tightly confined along one or more directions,

whose motion is governed by some lower dimensional forms of Equation 15. Suppose that the atoms are confined near the plane $z = 0$, and that their order parameter separates as $\psi(s, x, y, z) = \psi(s, x, y)\theta(z)$. This occurs when the potential confining the atoms along the z -axis is too tight for them to occupy excited orbitals, and θ is the ground orbital of that potential. In this case, the two-dimensional order parameter $\psi(s, x, y)$ satisfies a differential equation of the same form as Equation 15,

$$2i\partial_s\psi = \left(-\nabla_2^2 + K^2 + 8\pi a(\int |\theta|^4)|\psi|^2\right)\psi. \quad (16)$$

The quantity $\int |\theta|^4$ is the density-weighted mean density of the ground orbital. Its reciprocal is the length of a box in which an atom would be held with the same average density as in the orbital, which is a reasonable way to define the length of the orbital. In the special case that θ is the ground state of a harmonic oscillator with characteristic length l_0 , this length is $1/\int |\theta|^4 = \sqrt{2\pi}l_0 \approx 2.5l_0$. For this tight harmonic potential, Equation 16 reduces to

$$2i\partial_s\psi = \left(-\nabla_2^2 + K^2 + \sqrt{32\pi}\frac{a}{l_0}|\psi|^2\right)\psi \quad (17).$$

The factor a/l_0 is dimensionless, and the numerical factor is $\sqrt{32\pi} = 10.027$. In this confined condition, we will refer to the quantity $\xi^{-2} = \sqrt{32\pi}\frac{a}{l_0}|\psi|^2$ as the healing length. The three dimensional healing length is not a physically meaningful quantity—the atoms are packed so tightly along the xy -plane that there is no volume of condensate with that extent.

For general forms of θ , we can define an effective length $l_0 = (\sqrt{2\pi}\int |\theta|^4)^{-1}$, and Equation 17 remains valid. However, the physical assumption that the wave function along the z axis is fixed to θ is only valid when the atoms are tightly confined in one plane. When the condensate is split into several components along the z -axis, those components may have different states. In this case the ground orbital θ is degenerate.

Now suppose that the atoms are confined to a line around the x axis, with $\psi(s, x, y, z) = \psi(s, x)\theta(y, z)$. Equation 15 reduces to

$$2i\partial_s\psi = \left(-\partial_{xx}^2 + K^2 + 8\pi a(\int |\theta|^4)|\psi|^2\right)\psi. \quad (18)$$

In this case, $\int |\theta|^4$ is the reciprocal of the area of the orbital θ . When θ is the ground state of a plane harmonic oscillator with length scale l_0 , this is $\int |\theta|^4 = 1/(2\pi l_0^2)$, and Equation 18 reduces to

$$2i\partial_s\psi = \left(-\partial_{xx}^2 + K^2 + \frac{4a}{l_0^2}|\psi|^2\right)\psi.$$

In the classical wave limit, the atomic density does not matter, except through its effect on the healing length ξ . But in the regime of quantum optics that extends this classical limit, the discreteness of atoms sets the scale of quantum noise. It is convenient to make the dispersion unit a notional atomic separation $1/n$. This makes ψ a quantity of order one, and gives Equation 18 the dimensionless form

$$2i\partial_t\psi = \left(-\partial_{xx}^2 + K^2 + 2\gamma|\psi|^2\right)\psi. \quad (19)$$

The number $\gamma = 2a/l_0^2 n = 1/2n^2 \xi^2$ is called the Lieb-Liniger parameter. It is the squared ratio of the atomic separation $1/n$ to the healing length ξ . This measures the extent to which quantum mechanics is necessary to describe the condensate. Classical mechanics is adequate when there are many atoms per healing length, because classical phenomena involving the order parameter happen on the scale of the healing length and longer, while shot noise becomes significant on the scale of the atomic separation and shorter.

9. Graphs of wave functions

This thesis will consider many functions on phase space. These are functions of a complex variable whose real part is proportional to position, and imaginary part to momentum, according to Equation 4. A simple example is $f(\alpha) = \langle \alpha | \psi \rangle$, where $|\psi\rangle$ is a fixed state. It is notoriously difficult to draw graphs of complex functions, because the domain and range have two dimensions each, which does not fit into the three dimensional space we live in.

The approach used here, following [33], is to plot the values $f(\alpha)$ on the complex plane of the variable α . The brightness at each point is proportional to the modulus of $f(\alpha)$, and the colour of the point indicates the phase of $f(\alpha)$ on the scale shown in Figure 3. Red-brown represents positive real values, and blue-grey represents negative reals. The most saturated green and purple colours represent purely imaginary values.




Figure 3. The colour scale used to represent complex phase.

The images in Figure 4 show $\langle \alpha | \psi \rangle$ as a function of the complex amplitude α , where the state $|\psi\rangle$ is respectively the vacuum state, a coherent state and a number state. The identical size and shape of the blobs for the vacuum and coherent states reflects Equation 7. The stripes on the coherent state are caused by the factor of $\exp(i \text{Im } \alpha^* \beta)$ in Equation 6. Note that $\text{Im } \alpha^* \beta$ can be interpreted geometrically as a cross product $\alpha \times \beta$, where α and β are identified with vectors in the complex plane. Because of Equation 8, these phase space functions can be interpreted as expansions $|\psi\rangle = \int f(\alpha) |\alpha\rangle d\alpha$. Note that the coherent states are overcomplete, so these expansions are not unique. For example, any phase space function whose modulus has circular symmetry and whose phase is $\exp(in \arg(\alpha))$ is an expansion of the number state $|n\rangle$.

The colors on these plots do not convey any extra information about the state $|\psi\rangle$: Bargmann showed that the mapping $f(\alpha) = e^{\frac{1}{2}|\alpha|^2} \langle \psi | \alpha \rangle$ is analytic [34], so the phase of

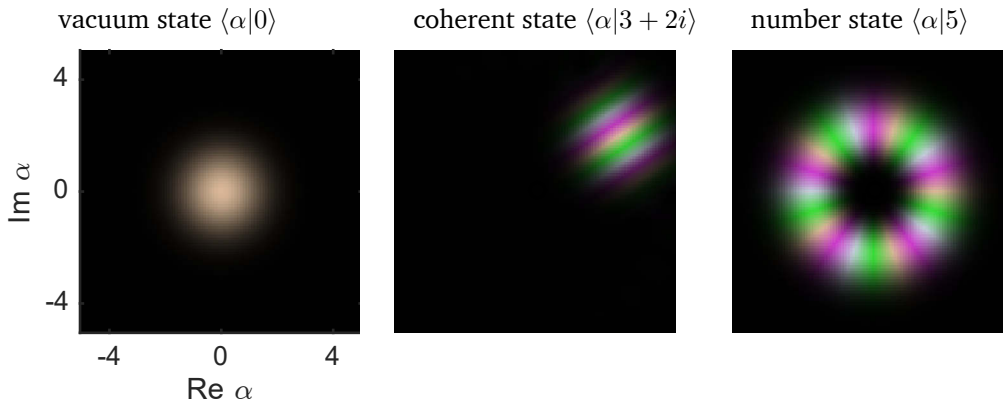


Figure 4. Examples of phase space coefficients.

$\langle \psi|\alpha\rangle$ is determined by its modulus. Equivalently, $|\langle \psi|\alpha\rangle|^2$ is the Q distribution, from which we could in principle reconstruct $|\psi\rangle$ and then the phase. (In practice, reconstructing $|\psi\rangle$ from its Q distribution is difficult, for reasons discussed in Section 16.) But a large part of this thesis discusses expansions of kets over coherent states; the phases usefully illustrate the function that this expansion needs to discretise. It is clear from the above plot that an expansion of the number state over coherent states would have to include enough of them to sample the corrugated phase without aliasing. This would not be apparent from a Q distribution for the number state, which is simply an annulus.

10. Condition and stability

All data are uncertain, and the quantities derived from them only represent reality approximately, to a finite precision. This is the case when we calculate the quantities by hand; it remains the case when we use a computer, and the more complicated methods that computers allow us employ can magnify the uncertainties in the input data in ways that are not obvious.

When a quantity is determined by experiment, there are usually several methods to go about it. Each method starts by making a different set of measurements, whose results determine the quantity in a specific way. Invariably, some of those methods will require very precise measurements to determine the quantity to a certain precision, and others will be more lenient. The different relations between each set of measurement results and the derived quantity cause errors to propagate in different ways.

In numerical analysis, there is a name for this type of relation between the precision of inputs and that of outputs: condition. Condition is a property of questions. The uncertainties in the measurements set the precision with which the question is posed, and the mathematical relation between the measurement results and the derived quantity determines the precision with which the answer is determined.

In simple data reductions, it is usually obvious how to calculate the desired quantity to the precision with which it is determined by the data. But in more complicated numerical

methods, measurements are not the only source of uncertainty. There are rounding errors in intermediate values, and discretisation errors when continuous functions are sampled on grids. These things can cause the calculated quantity to be uncertain even when the measured data actually determine it precisely. So as well as asking how precisely the question determines the answer, it is also necessary to consider whether the algorithm used to calculate the answer achieves that precision.

This property of numerical methods is called stability. Condition is a property of questions, but stability is a property of the algorithms that calculate the answers. A stable algorithm calculates the answer as precisely as the input data determine it. Usually, a numerical algorithm decomposes the problem into a series of steps. In a stable algorithm, each step is well-conditioned, and unstable algorithms result some of the steps being unnecessarily badly conditioned.

Some examples might help illustrate this. The canonical example of a badly conditioned problem is to calculate $\sin(x)$ for a value of x near 10^{10} . A relative uncertainty in x of $10^{-10} \times \pi$ causes a 100% uncertainty in its sine. This is by the nature of the sine function, and remains the case no matter how its value is computed.

The canonical example of a well-conditioned question, where the obvious way to answer it is unstable, is computing $1 - \cos(x)$ for $x = 10^{-10}$. This is very easy to do: the answer is 5×10^{-21} to around 10 digits. However, the obvious way to do it with a computer would be to determine the value of $\cos(x)$ numerically, and subtract this from 1. Unless the computer stores numbers to 21 digits, the result will be a random number whose order of magnitude is the precision of the computer.

When experiments are done on many-particle systems, phenomena such as interference and superconductivity are reproducible, and changing the experimental conditions slightly does not greatly change the phenomena that result. This suggests that the properties of these systems are well-conditioned functions of the experimental parameters. However, methods of simulating these phenomena on computers have tended to be very unstable. The challenge is to find stable ways of computing these results, or rather the values predicted for them by quantum mechanics. Chapter 5 will consider this at length.

In numerical analysis, just like in experimental uncertainty analysis, we consider the uncertainties to be small. This allows their propagation to be analysed by linear approximations. If a vector of results y is determined by a vector of data x as $y = f(x)$, then we can suppose that the uncertainties relate as $\Delta y = f'(x) \cdot \Delta x$, although perhaps the expected values of the components of Δx will add in quadrature. The steps in numerical algorithms are generally linear problems, so much uncertainty analysis reduces to the condition of some linear problem.

This is done with a matrix factorisation called the singular value decomposition. This produces singular values and singular vectors, which are similar to eigenvalues and eigenvectors. However, the domain and range of f might be different spaces, so the Jacobian matrix f' might not be rectangular. There are two different types of singular vectors, one for the domain and one for the range.

For a matrix A , the singular value decomposition is $A = USV^\dagger$, where U and V are orthogonal square matrices, and S is a diagonal matrix. This exists and is unique for every matrix [35]. The orthonormal basis defined by the columns of V is called the right singular vectors of A , and the columns of U are the left singular vectors. When we write $A: R \rightarrow T$, where R is the domain and T is the range, the left singular vectors lie in T and the right singular vectors in R . The diagonal elements of S are real, and they are called the singular values.

Geometrically, the factor V^\dagger chooses a set of axes, the diagonal matrix S scales each axis by its singular value, and the matrix U rotates the axes into the domain space. The geometric meaning is that every linear transformation comprises choosing axes, scaling the axes, and rotating the axes.

Rotating axes does not affect the condition of A , so it depends on the singular values only. More precisely, it depends on their spread. Suppose one singular value is small, and the input data lie parallel to the corresponding singular vector. If these data are affected by an error that lies along a singular vector with a large singular value, then the relative change in Ax will be large. Therefore the condition of a linear transformation is measured by the ratio of its largest and smallest singular values, which is called its condition number. A value of 1 is the best condition possible, where the relative error in the result is the same as in the data. Larger values indicate worse condition, and the number diverges to infinity as A approaches a singular matrix.

11. Stiffness in quantum dynamics

Quantum dynamics is continuous. The quantities of physical interest are the expectation values of observables, of the form $\langle A \rangle = \langle \psi | A | \psi \rangle$, where A is a Hermitian operator; the state of motion $|\psi\rangle$ is a continuous function of time that satisfies the Schrödinger equation. The context for numerical quantum dynamics, and the extent to which it can possibly succeed, are set by the condition of these problems. If the state $|\psi\rangle$ is perturbed slightly, replacing it with a state $|\psi\rangle + |\delta\psi\rangle$, does this cause a small or a large change in the expectation values, and in the trajectory that the system follows from $|\psi\rangle$ as an initial state?

First consider the condition of the mapping that takes $|\psi\rangle$ to $\langle A \rangle$, where A is a fixed observable. For a small change in $|\psi\rangle$ of a fixed size, the largest relative change in $\langle A \rangle$ occurs when $|\psi\rangle$ is an eigenvector of A belonging to the smallest eigenvalue, and $|\delta\psi\rangle$ is an eigenvector belonging to the largest eigenvalue. This mapping from states to expectation values will be badly conditioned when the observable has some eigenvalues that are much larger than its expectation value; that is something we have to live with.

In Chapter 3, we will look at phase space Monte-Carlo methods for simulating quantum dynamics. These represent $|\psi\rangle$ implicitly, as a set of samples that are directly related to observable values. The virtue of these representations is that the expectation values are guaranteed to be well conditioned functions of the amplitudes in the ensemble. However, this is also a problem: when we want to add a perturbation $|\delta\psi\rangle$ whose expectation value is very different from that of $|\psi\rangle$, the only way to do it is to shift the sampled amplitudes a

long way. The implications of this for stability are discussed in Section 39.

The dynamics themselves are perfectly conditioned as a function of the initial state. The mapping $|\psi_0\rangle \rightarrow |\psi(t)\rangle$ is linear and unitary, so its condition number is one. The condition of $|\psi(t)\rangle$ as a function of t is more complicated. The condition of $e^{-iHt/\hbar}|\psi_0\rangle$ depends on the eigenvalues of the Hamiltonian. In particular, it depends on their spread: by adjusting the zero of potential energy, the eigenvalues can always be made to lie in an interval around zero without changing the dynamics.

The coefficients of the energy eigenvectors can be visualised as samples of a complex exponential $e^{-iE_it/\hbar}|\psi_0\rangle$, which becomes more tightly twisted as time increases. If the range of the eigenvalues is ω , and they are roughly evenly spread, then the state will twist to become orthogonal to itself in a time around ω^{-1} . Quantum states that include components with a large range of energies have the same sort of conditioning problems that all oscillating functions do at large times.

In order to solve quantum dynamics numerically, the continuous states and times must be discretised to a finite set of numbers. The algorithms by which this is done have various levels of stability. This is almost always done by a method of lines, meaning that the state is represented by a set of parameters, and these parameters satisfy a system of ordinary differential equations [3]. The parameters could be samples of a wave function, or more general variational parameters.

Given that a method of lines is being used, the stability of the dynamics is determined by two things: the condition of the ordinary differential equations for the parameters, and the stability of the finite difference formula that calculates their solution. The condition of the equations is determined by the way the quantum state is discretised. This does not appear to have been studied very much, and Chapter 5 will discuss the condition of a certain type of discretisation.

On the other hand, the stability of finite difference formulae for ordinary differential equations has been studied very extensively; this is one of the canonical topics of numerical analysis. It is fairly straightforward to derive finite difference formulae that converge to the continuous solution in the limit of small steps. The problem is a phenomenon called stiffness. Many of these formulae diverge at moderate steps, before they finally converge to unnecessarily precise values at impractically small steps.

The standard form for a system of ordinary differential equations composes the variables into a vector-valued function z , so that the equations have the form

$$z'(t) = \zeta(z(t)), \quad (20)$$

where z' is the derivative of z . The usual formulae for solving this have the form [36]

$$\sum_{i=0}^N a_i z_{j+i} = h \sum_{i=0}^N b_i \dot{z}_{j+i}. \quad (21)$$

Time has been discretised at points t_j with step $h = t_{j+1} - t_j$. The sequence of vectors z_j satisfies the discrete difference formula of Equation 21, and z_j is intended to approximate $z(t_j)$.

The sequence \dot{z}_j is defined by $\dot{z}_j = \zeta(z_j)$. Unlike the function z' , this can be computed without knowing the exact solution z . This formula determines the value z_{j+N} in terms of z_j through z_{j+N-1} , and the derivatives $\zeta(z_i)$ evaluated at those points.

Consider the j th time step. The sequence z_i can be transformed into a new sequence Z_i with the substitution $z_i = z_j + (i-j)h\dot{z}_j + Z_{i-j}$. With this substitution, Equation 21 becomes

$$\sum_{i=0}^N a_i Z_i = h \sum_{i=0}^N b_i \dot{Z}_i, \quad (22)$$

where $\dot{Z}_i = \dot{z}_{j+i} - \dot{z}_j = \bar{\zeta}(Z_i)$. The conditions $\sum_{i=0}^N a_i = 0$ and $\sum_{i=0}^N i a_i = \sum_{i=0}^N b_i$ follow from the need for z_i to increase arithmetically when \dot{z}_i is constant, and the function $\bar{\zeta}(Z) = \zeta(Z + z_j) - \dot{z}_j$ satisfies $\bar{\zeta}'(Z) = \zeta'(Z + z_j)$. This substitution sets the step number to zero and sets $Z_0 = \dot{Z}_0 = 0$, but preserves the form of Equation 21 and the Jacobian of ζ .

If we suppose that h is small enough that the differential approximation $\zeta(Z_i) = \bar{\zeta}'(0) \cdot Z_i = \zeta'(z_j) \cdot Z_i$ holds for Z_0 through Z_N , then Equation 22 becomes

$$\sum_{i=0}^N (a_i - h b_i \zeta'(z_j)) Z_i = 0. \quad (23)$$

This can clearly be satisfied if the Z_i are multiples of an eigenvector w of the Jacobian $\zeta'(z_j)$, say with $\zeta'(z_j) \cdot w = \lambda w$. Moreover, because the Z_i are all generated by Equation 22, they can grow steadily, with $Z_i = c^i w$, provided that c satisfies the polynomial

$$\sum_{i=0}^N (a_i - h \lambda b_i) c^i = 0.$$

The above analysis applies at every time step, so the question of stability is now simple. If $|c| > 1$, any component of w will grow exponentially in the sequence z_i . When the z_i are computed with finite precision, it is almost inevitable that roundoff will introduce such a component, and the series z_i will diverge exponentially from the solution $z(t_i)$ if there is any eigenvalue λ for which $|c| > 1$. The solution will be stable if $|c| \leq 1$; if we want to solve the differential equation, instead of generating a sequence of numbers that shrink to zero, we are usually stuck on the boundary with $|c| = 1$. For a given formula, with fixed a_i and b_i , one can draw a region of stability in the complex plane, such that the formula is stable when the quantity $h\lambda$ lies within the region and unstable otherwise. Inevitably, this region is bounded, and large eigenvalues from directions where $\zeta(z)$ varies sharply with z require a small time step h for the formula to be stable.

When Schrödinger's equation is solved by a method of lines, the resulting ordinary differential equations tend to be stiff, because the Hamiltonian includes a kinetic energy. The eigenvalues of kinetic energy are squared wave numbers. When the wave function is sampled on a grid, the spacing of grid points determines the largest wave numbers that can be represented by the samples, and therefore the largest eigenvalues of the discretised kinetic energy. Schrödinger's equation is linear; when it is discretised on a grid and put in

the form of Equation 20, the function ζ is a linear operator on the vector z , corresponding to a discretised Hamiltonian. The Jacobian of a linear operator is the same operator. Therefore the largest eigenvalues, which constrain the step at which the the formulae become stable, are simply the largest energies of states whose wave functions can be represented on the grid, and these increase as the grid spacing shrinks. This is a general problem, which causes many PDE formulae to become unstable when the time step is disproportionately large compared to the grid spacing.

This problem of stiffness due to large energies is not avoidable if we want to simulate phenomena that occur at those energies. However, as we will discuss below, it is possible to make the ordinary differential equations much stiffer than they need to be, and precautions must be taken to avoid that.

On the other hand, if the high energies are not part of the simulated physics, and only occur because of a fine grid spacing, the growth of parasitic modes can be prevented. The stiffness of the GPE comes from rapid phase rotations. Formulae that look like $\psi_{n+1} = \psi_n + h\psi'_n$ fail to follow these, as is clear from an Argand diagram:




Figure 5. Eigenmodes with large frequencies become unstable when time is discretised.

One solution is also apparent. Instead of the approximation $\psi(s+h) = \psi(s) + ih\mathcal{E}\psi(s)$, we should try to approximate $\psi(s+h) = e^{ih\mathcal{E}}\psi(s)$, in such a way that the energetic components will rotate instead of growing. The operator \mathcal{E} splits into two terms,

$$\mathcal{E}_k = -\frac{\nabla^2}{2} \quad \mathcal{E}_x = \frac{K^2}{2} + 4\pi a|\psi|^2.$$

The operator $e^{i\mathcal{E}_x}$ can be applied very simply to a function $\psi(x)$, while $e^{i\mathcal{E}_k}$ can be applied simply given the Fourier transform $\psi(k)$.

12. Frames

When a continuous object is reduced to a finite set of parameters z_i , a very common way to do so is by taking components over a basis. The preceding section discussed taking samples of a wave function on a grid of points. Later in this section, we will discuss the equivalence between this and expanding the function over a basis of sinc functions. Later in the thesis, we will need to consider more general expansions, over sets that are not quite bases but perform the same numerical function. These sets are called frames.

A basis is a set of vectors that span a space, and that are linearly independent. The first property ensures that every vector in the space can be expanded as the sum of its components along the basis. This property is obviously useful when the components are used for

calculations. The second property, linear independence, ensures that the components are unique. This is very useful when bases are used for abstract purposes: for example, it allows vectors to be counted without worrying that some of them might be counted multiple times as different sets of components for the same vector. But it is less obviously important for calculations. Provided we have some way to represent a vector on a computer, it rarely matters that there are other ways that it could be represented.

In numerical work, we will often identify the elements of a basis $\{a_1, \dots, a_n\}$ with column vectors, and compose them into an expansion matrix $A = (a_1 \dots a_n)$, so that any vector x can be expanded in the form $x = Ac$ where c is a column vector of coefficients. Linear independence does ensure that this matrix is invertible. This guarantees that the matrix has a positive and finite condition number, but says nothing about its size. The components over a basis are unique in principle, but when the components of two vectors that are important in a calculation differ by less than the round off error, the matrix is effectively singular for calculation purposes.

The basis property can also be quite fussy; it excludes some methods of representation that don't seem to cause any real problems. Consider a set of coherent states whose amplitudes form a grid in phase space. This is an obvious set to expand kets over, and, provided the grid is dense enough, these expansions are not problematic. However, it can be shown that these kets form a basis only under a very specific condition—the grid density must be such that there is one amplitude per area π of phase space, and one point on the grid must be excluded [37]. Grids more sparse than this fail to span phase space, and denser grids are not linearly independent.

For these reasons, numerical analysts have generalised the concept of a basis to that of a frame. The elements of a frame span a space, but they need not be linearly independent: instead, they are required to compose an expansion matrix whose condition number is finite and nonzero. Mathematicians have traditionally stated this in a more obscure way, that generalises neatly when the set of vectors is continuously infinite. This is described below, after we give an example.

For frames with one element for each dimension of the vector space, whose expansion matrices are square, a nonzero condition number ensures that the matrix is invertible; in this case a frame is the same thing as a basis. When a basis is defined in terms of condition number, it is possible to say that sets of vectors with small condition numbers are good bases, whereas those with large condition numbers are almost linearly dependent. We will refer to these as snug frames and loose frames. When the condition number is one, a basis is orthonormal. The generalisation of this to rectangular expansion matrices is called a tight frame.

To introduce the idea, consider a frame that spans \mathbb{R}^2 , the simplest vector space possible, illustrated in Figure 6. The vector a_1 has length a along the horizontal axis, while a_2 is a unit vector at an angle θ . There is a degenerate case where θ is an integer multiple of π and both vectors are horizontal, and another where a shrinks to zero. Apart from these cases, the frame is a basis for \mathbb{R}^2 .




Figure 6. Example of frame bounds for a basis.

The expansion matrix for this frame is

$$A = \begin{pmatrix} a & \sin \theta \\ 0 & \cos \theta \end{pmatrix}.$$

It follows from a simple calculation that for a unit vector c ,

$$\frac{d - \sqrt{d^2 - 4 \sin^2 \theta}}{2a} \leq \|Ac\|^2 \leq \frac{d + \sqrt{d^2 - 4 \sin^2 \theta}}{2a},$$

where $d = (a + \frac{1}{a})/|\sin \theta|$, and that these bounds are achieved. The condition number of A is the ratio of the bounds on $\|Ac\|$,

$$(\text{cond}A)^2 = \frac{d^2}{2} + d\sqrt{\frac{d^2}{4} - 1} - 1. \quad (24)$$

The limit of bad condition occurs when d is large, and in this case $\text{cond}A \approx d$.

There are two cases where the condition number of Equation 24 becomes large. In these cases, the ratio between the norm of $x = Ac$ and the norm of c can be large or small. They are obvious from geometry.

The first case is where θ is small. The point $\theta = 0$ is the degeneracy where a_1 and a_2 are parallel; the condition number becomes infinite in this limit because $|\sin \theta|$ decreases to 0. When θ is small but not zero, the vectors are a basis. However, some vectors are almost orthogonal to both basis elements, and are expanded as the difference of two large and nearly equal vectors. This is something that one always wants to avoid in numerical work; in variational dynamics, the reasons to avoid it are stronger than usual, because it will result in stiff differential equations for the variational parameters.

The other case occurs where a is small or large, so that a_1 and a_2 have very different lengths, and either a or $\frac{1}{a}$ is large. Again, the poor condition of A is geometrically obvious. When x is parallel to the large vector, $|c|$ will be small compared to $\|x\|$, and vice-versa when x is parallel to the small vector. A combination of the two cases, where the vectors are almost parallel and have very different lengths, results in d being the product of the two factors, and a multiplicatively worse condition number.

Mathematicians traditionally define a frame in terms of frame bounds. These are positive numbers L and U such that

$$L\|x\|^2 \leq \sum |\langle x, a_i \rangle|^2 \leq U\|x\|^2. \quad (25)$$

Note that the middle term is also $|A^\dagger x|^2$. The condition number of A^\dagger is the reciprocal of the condition number of A so, if the lower bound L and the upper bound U are achieved, their ratio is the condition number of A . If the frame vectors a_i are all multiplied by a scalar, then both bounds are multiplied by the same scalar, so only the ratio is really significant. However, it can be useful to distinguish the problem where L becomes small, and the frame is almost linearly dependent, from the problem where U becomes large, and the frame barely spans the space.

The expansion matrices of frames can also be rectangular, when the frame has more elements than the space has dimensions. This is most interesting in the case of infinite dimensions, where for instance the frame might be a grid of coherent states with too great a density to be a basis; this is discussed in the next section. Such frames can still be tight, with a condition number of one. These tight frames share many of the useful properties of an orthonormal basis.

An example of such a redundant frame is shown in Figure 7. The elements are three unit vectors at angles 0 , $\pi - \theta$ and $\pi + \theta$, with $0 \leq \theta \leq \frac{\pi}{2}$. These three vectors lie in a two dimensional space, and they are not linearly independent. With this frame, the expanded vectors in the plane have two elements, and the coefficient vectors have three, so the expansion operator A is a 2×3 rectangular matrix

$$A = (a_1 \ a_2 \ a_3) = \begin{pmatrix} 1 & -\cos \theta & -\cos \theta \\ 0 & \sin \theta & -\sin \theta \end{pmatrix}.$$

This expansion operator is not invertible, because it is a mapping between two spaces of different dimensions. It is onto except in the degenerate case $\theta = 0$, but it is never one-to-one.




Figure 7. Example of a redundant frame.

The singular value decomposition of A is

$$A = USV^\dagger = (u_1 \ u_2) \begin{pmatrix} \sigma_1 & 0 & 0 \\ 0 & \sigma_2 & 0 \end{pmatrix} (v_1 \ v_2 \ v_3)^\dagger.$$

Although A is not invertible, it still has a finite condition number $\text{cond} A = \frac{\sigma_1}{\sigma_2}$, except in the degenerate case. In this form, the lack of one-to-one is explicit. We can write the product SV^\dagger as the sum of two block products,

$$A = U \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} (v_1 \ v_2)^\dagger + U (0 \ 0) v_3^\dagger.$$

The first term multiplies the components of v_1 and v_2 by an invertible matrix, but the second term multiplies v_3 by the bottom row of S , which is zero. This suggests a way to define an operator called the pseudoinverse of A ,

$$A^+ = (v_1 \quad v_2) \begin{pmatrix} \frac{1}{\sigma_1} & 0 \\ 0 & \frac{1}{\sigma_2} \end{pmatrix} U^\dagger = V \begin{pmatrix} \frac{1}{\sigma_1} & 0 \\ 0 & \frac{1}{\sigma_2} \\ 0 & 0 \end{pmatrix} U^\dagger.$$

This inverts A in the sense that $A^+A = 1$. If $y = Ax$, then $x = A^+y + cv_3$, where cv_3 is a vector of coefficients along the direction that A annihilates. The right singular vectors v_i are an orthonormal basis, so the preimage x has the least norm when $c = 0$, and $x = A^+y$. The matrix AA^+ projects coefficient vectors onto these least norm expansions.

The coefficients of an expansion over a frame are not necessarily unique: adding a component of v_3 to an expansion does not change the expanded vector. However, the expansion A^+v is unique in that it is orthogonal to the null space of A , and therefore it has a smaller norm than any other expansion.

The right singular vector v_3 , corresponding to the zero singular value, lies in the direction $(\cos\theta, \frac{1}{2}, \frac{1}{2})$. From symmetry, we can guess that the left singular vectors are the axes of \mathbb{R}^2 . Let e_1 be a horizontal unit vector, and e_2 a vertical one. The left singular vector v_2 that maps to e_2 is in the direction $(0, 1, -1)$, and has singular value $\sigma_2 = \sqrt{2} \sin\theta$. The third singular vector v_1 maps to e_1 , so it is in the direction $(1, -c, -c)$ for some value of c . Given that it must be orthogonal to v_3 , the direction is $(1, -\cos\theta, -\cos\theta)$, and the singular value is $\sigma_1 = \sqrt{1 + 2\cos^2\theta}$. These two singular values are the frame bounds.

For $\theta = \frac{\pi}{3}$, the directions of the frame vectors are evenly spread. At this point, the singular values are equal, the condition number is 1, and the frame is tight. In the degenerate case $\theta = 0$, the singular value σ_y drops to zero, and these vectors cease to be a frame. For $\theta < \frac{\pi}{3}$, all three vectors are nearly horizontal. In this case, $\sigma_1 > \sigma_2$, and the condition number is $\sigma_1/\sigma_2 = \sqrt{\frac{3}{2\sin^2\theta} - 1}$. For $\frac{\pi}{3} < \theta \leq \frac{\pi}{2}$, $\sigma_2 > \sigma_1$ and the condition number is the reciprocal of that.

For $\theta = \frac{\pi}{2}$, two of the vectors are degenerate, but the condition number is only $\sqrt{2}$. This expansion matrix is not particularly badly conditioned.

13. Fourier series as frames

The frames in the last section had a finite number of elements; this allowed us to draw them as vectors on a page. However, it is also possible to have frames with continuously infinite numbers of elements. When the goal is to represent continuous mathematical objects, such as functions or kets, it is often useful to consider such frames. In the discrete case, a linear combination is a sum $b = \sum a_i c_i$, but continuous objects are more usually composed from infinitesimal linear combinations of the form $b = \int a(k)c(k) dk$.

In the numerical case, the number of frame kets is always finite; unless $|A\rangle$ is singular, the frame bounds always exist, with A and B on the order of the largest and smallest singular values of $|A\rangle$. The mathematical definition of a frame needs to be loosened somewhat

for numerical purposes. A bound of the order 10^{15} might as well be infinite, and one of the order 10^{-15} might as well be zero. Also, a finite set of states can not be a frame for the whole of Hilbert space, so we need to think about what space $|\psi\rangle$ can lie in. There is a question about how accurately and stably elements just outside the space spanned by a frame can be approximated by expansions over that frame, and this question still appears to be open.

The most familiar continuous frame is the Fourier transform, where the vector is a function $b(x)$, and the frame elements are complex exponential functions $a(k) = e^{ikx}$. Frames were first introduced in the context of nonharmonic Fourier series [38], expansions over complex exponentials whose frequencies are not commensurate; the discussion in this thesis is limited to ordinary Fourier series. We will approach these from the viewpoint of functions of a limited bandwidth W . These can be specified by their values on a dense enough grid, and a function can be recovered by interpolating these samples with sinc wave packets.

There are many ways to normalise the Fourier transform and its inverse. We will use the one where

$$f(x) = \frac{1}{\sqrt{2\pi}} \int F(k) e^{ikx} dk \quad \text{and} \quad F(k) = \frac{1}{\sqrt{2\pi}} \int f(x) e^{-ikx} dx. \quad (26)$$

This has the advantage that

$$\int F^*(k) G(k) dk = \int f^*(x) g(x) dx,$$

so that the Fourier transform is unitary under the \mathcal{L}^2 inner product.

According to Equation 5, the wave functions of coherent states are Gaussian wave packets. Graphs of these look similar to graphs of sinc wave packets, so we might guess that expanding a ket over coherent states is similar in some ways to the more familiar process of interpolating samples with sinc wave packets. When we discuss expansions over coherent states, in Section 40, this will indeed turn out to be the case.

Numerical frames are discrete because computer storage is finite. However, consider a continuous frame with elements $a(k)$, where k varies over the real numbers. A discrete frame with many elements $a(k_n)$, whose adjacent elements are nearly parallel vectors, will have linear combinations $b = \sum a(k_n) c_n$ that approximate linear combinations $b = \int a(k) c(k) dk$ in the continuous frame. In this case, many properties of the discrete frame are determined by properties of the continuous frame, and the continuous frame can be much easier to analyse.

Suppose that a function f has bandwidth W . This means that its Fourier transform F fits into the interval $C = [-W, W]$ on the reciprocal axis. The function F can be expanded as a Fourier series

$$F(k) = \frac{\chi_C(k)}{\sqrt{2W}} \sum_n F_n e^{n\pi ik/W}, \quad (27)$$

where the hat function χ_C takes the value 1 on C and 0 elsewhere. The Fourier series is periodic, and this selects one copy of F instead of repeating it along the reciprocal axis.

The function $\chi_C/\sqrt{2W}$ is the Fourier transform of a sinc wave packet

$$a(x) = \sqrt{\frac{W}{\pi}} \frac{\sin Wx}{Wx} = \sqrt{\frac{W}{\pi}} \operatorname{sinc}(Wx).$$

Modulating χ_C in reciprocal space with $e^{\pi ik/W}$ has the effect of translating a by $-\pi/W$. Therefore, Equation 27 is the Fourier transform of an expansion of f over a grid of sinc wave packets,

$$f(x) = \sum_n F_n a\left(x + \frac{n\pi}{W}\right).$$

The series terms $(2W)^{-\frac{1}{2}} \chi_C(k) e^{\pi ik/W}$ are orthonormal, and the Fourier transform is unitary, so the wave packets $a\left(x + \frac{n\pi}{W}\right)$ form an orthonormal basis for the space of functions of bandwidth W . The grid density W/π is called the Nyquist density. When the bandwidth is expressed as a circular frequency $\nu = W/2\pi$, this becomes the familiar result that the Nyquist density is twice the bandwidth.

The coefficient of $e^{n\pi ik/W}$ in the Fourier series 27 is

$$F_n = \frac{1}{\sqrt{2W}} \int_{-W}^W F(k) e^{n\pi ik/W} dk.$$

The Fourier transform F is zero outside the bounds of the integral, so we can extend them to give

$$F_n = \frac{1}{\sqrt{2W}} \int_{-\infty}^{\infty} F(k) e^{n\pi ik/W} dk = f\left(-\frac{n\pi}{W}\right).$$

But the value F_n is also the coefficient of the sinc wave packet centred at $-\frac{n\pi}{W}$. In other words, the coefficients of the wave packets are samples of f .

We will now consider grids with different spacings, and wave packets of different widths; not all of these sets of functions span the space. Suppose that grid has a density K/π other than the Nyquist density, and the wave packets are $a_n = \operatorname{sinc}\left(x + \frac{n\pi}{K}\right)$. This density is the period of the Fourier series in reciprocal space, so in order to expand f over these wave packets, it would be necessary to expand F on the interval $[-W, W]$ using a Fourier series with period $2K$,

$$F(k) = \frac{\chi_C(k)}{\sqrt{2W}} \sum_n F_n e^{n\pi ik/K}. \quad (28)$$

This series can only reproduce F when $K \geq W$. The copy of F displaced by K along the reciprocal axis has support at wave numbers above $K - W$, and this overlaps the original F when $K < W$. As is well known, an attempt to reproduce f by interpolating on an undersampled grid produces an aliased function instead of f .

When the grid is oversampled, with $K > W$, sampling and interpolation still work. The interval $[-W, W]$ on which F is nonzero fits inside the period $[-K, K]$, as shown in Figure 8. The terms $\chi_C e^{n\pi ik/K}$ are no longer orthogonal, but all the other arguments made at the Nyquist density are still valid, and the oversampled wave packets span the space of functions of bandwidth W . However, these wave packets are linearly dependent.




Figure 8. Oversampled sinc wave packets in reciprocal space.

A Fourier series $\sum_n F_n e^{n\pi ik/K}$, whose coefficients F_n are not all zero, can be zero on the interval $[-W, W]$ but nonzero on the interval $[W, K]$. When such a series occurs in Equation 28, the right hand side is the Fourier transform of a linear combination of sinc wave packets with nonzero coefficients F_n , but this linear combination sums to the zero function.

This set of oversampled sinc functions forms a frame for the space of functions of bandwidth W . This can be useful. A sinc wave packet decays slowly away from the origin, proportional to $1/|x|$. This is necessarily the case in order for its Fourier transform to cut off sharply at the wave number K , but it causes a function interpolated using these wave packets to depend on samples at distant points. When $K > W$, we can add to the sinc wave packet a function g whose Fourier transform G decays smoothly from $G(W) = 1$ to $G(K) = 0$ on the interval $[W, K]$ and is zero elsewhere, with the result that the wave packets $a + g$ decay faster than $1/|x|$. The Fourier transform F is zero on the interval $[W, K]$, so it remains the case that

$$F(k) = (\chi_C(k) + g(k)) \sum_n F_n e^{n\pi ik/K},$$

and the function f is still reproduced when its samples are interpolated by the narrower wave packets $a + g$. The slow decay of sinc wave packets is a consequence of pushing the sampling density right to its limit; results of this type, where the frame elements become singular in some limit where the set ceases to be a frame, are common when continuous frames are discretised.

Now, as well as changing the grid density, consider changing the width of the sinc functions, so that $a = \sqrt{\frac{J}{\pi}} \text{sinc}(Jx)$. We will call J the interpolation density. Interpolating with these wave packets constructs a function whose Fourier transform is

$$F(k) = \chi_{[-J, J]} \sum_n F_n e^{n\pi ik/K}.$$

This reproduces f provided that $W < J < 2K - W$. Aliasing occurs when the interpolation density is larger than this limit, and frequencies are truncated when it is less than the bandwidth, as in Figure 9.

The sinc wave packets are orthonormal when their width is matched to the grid spacing, with $J = K$. They are a basis for functions of bandwidth W when $J = K = W$. Narrow packets, with $J > W$, have bandwidth J and lie outside the space of functions with




Figure 9. The effect of the frequency of the wave packets.

bandwidth W . Their projections onto that space are the wave packets with $J = W$. Broad wave packets, with $J < W$, do not span the space.

Chapter 3 Phase space Monte-Carlo methods

Chapter 4 will use the truncated Wigner method to calculate the dynamics of some highly elongated atom traps. This chapter will derive and justify the method. Section 15 introduces an isomorphism between Hilbert space operators and phase space distributions, functions taking classical states to real numbers; provided that a density operator corresponds to a positive function, it can be identified with a probability distribution over phase space from which states can be drawn. That section will also describe how the expectation values of quantum observables can be calculated as the average of the values that its corresponding function takes at these sample states. Section 19 will justify why the quantum motion of the density operator can be identified with the classical motion of the sampled points in phase space, and examine the assumptions on which that identification relies.

This chapter will also describe ways to generalise the truncated Wigner method, introducing a family of related phase-space methods due to Cahill and Glauber [39, 40]. These variants trade off the range of states that correspond to positive functions against the convergence and accuracy of the sampled dynamics. They are introduced in Sections 16 and 18, then Section 19 discusses the Monte Carlo dynamics of the whole family.

Almost all of the results described in this chapter are derived in the papers cited. A few minor observations are not stated explicitly in any of those papers, but are obvious enough that they have very likely been published somewhere by now.

These phase space methods are one of only a handful of ways that the dynamics of a quantum field can be computed. For reasons that were elaborated in Section 2, it is not practical to represent the quantum state of an arbitrary field; any such method needs to propagate through time an incomplete specification of the field state, and to derive from it the expectation values of the interesting observables.

The other common method uses a perturbation theory for Green's functions [41, 4]. It either assumes that the system is a small perturbation of one whose dynamics are known exactly, or uses other approximations to sum the higher order terms in the perturbation series. Instead of this perturbative approximation, the phase space methods rely on a fair sampling approximation. It would require a very large ensemble of classical states to reconstruct the density operator, but hopefully the averages will converge to the interesting expectation values with a smaller ensemble that is practical to generate numerically. In both methods, the assumptions about convergence are questionable; but some assumption must be made, because the amount of information required to represent the state of many

quantum particles increases as the exponential of the number of orbitals they occupy.

All of these methods are based on correspondences between quantum states and distributions. When the state is represented by a density operator, these correspondences are linear mappings. Therefore, linear operations on the density operator are equivalent to linear operations on the corresponding distribution. In order for the state to obey the laws of motion, the distribution must satisfy a certain partial differential equation, which ensures that the density operator satisfies the quantum Liouville equation.

When phase space methods are applied to a quantum field, the calculations are nearly always done numerically. The field is discretised by expanding it over a set of modes; when the field is a cold gas, these are atomic orbitals. The dimension of the phase space is the number of orbitals, and the only practical way to discretise the distribution functions is to draw a random ensemble of points in the phase space. There is a special case where the distribution function is a probability distribution over phase space, and it satisfies a differential equation with a Fokker-Planck form. In this case, there is a way to derive a system of ordinary or stochastic differential equations such that initial conditions sampled from the distribution function will produce trajectories that continue to be distributed according to the distribution functions at later times. This allows the dynamics to be computed by a Monte-Carlo method [42].

Section 11 discussed how partial differential equations can be discretised by a method of lines, where a fixed grid of points is chosen in space, and the values of the function at these points are computed. The Monte-Carlo methods work in the opposite way. They do not store any values of the function explicitly, but encode them in the density of the sample points. Instead of fixing the points on a grid and computing the values by solving a system of differential equations, the trajectories of the points satisfy a set of differential equations while the values are stored only implicitly. This is a huge advantage in high dimensional spaces, because the number of points in a grid grows exponentially with the dimension; the ket space for a quantum field has dimension proportional to the resolution with which the field is represented, so a Monte-Carlo method is the only practical approach.

In all of these methods, the trajectory of each sample point is independent of the others, and satisfies its own differential equation. When the truncated Wigner method is applied to a closed system, this equation is the law of motion for a classical system whose Hamiltonian is the Wigner transform of the quantum Hamiltonian [24]. In the other methods from the family, it is not, and in general it is a stochastic law of motion. This independence allows efficient simulations with large ensembles, because the trajectories can be computed concurrently. Averaging moments over ensembles is an associative operation; so phase-space simulations are embarrassingly parallel, and it is practical to run them on very large scales.

14. Quartic oscillator

We will use the quartic oscillator as a test system for the simulation methods. This is similar to the harmonic oscillator, in that the Hamiltonian commutes with the number oper-

ator $a^\dagger a$. However, instead of the quadratic Hamiltonian $H = a^\dagger a$, there is a quartic one $H = a^{2\dagger} a^2 = n(n-1)$. It is difficult to picture this in first quantisation, in terms of forces and potentials acting on a particle; the momentum does not just enter into the Hamiltonian as $p^2/2m$, but there are terms proportional to p^4 . On the other hand, this is a very simple and common Hamiltonian for particles in second quantisation, where it represents particles that are confined to one orbital and that repel each other.

This is a good test system for a couple of reasons. Firstly, the quantum and classical dynamics are very different. The classical motion is very much like a harmonic oscillator, except that the frequency of the oscillations increases with their amplitude. The Hamiltonian is nonlinear, but the dynamics can not be chaotic because the phase space is only two dimensional. The trajectories of the system still describe circles in phase space, but the period to go around the circle changes with its size.

The quantum dynamics are shown in Figure 10. The system begins in a coherent state, the closest quantum analog to a classical oscillator. In a harmonic oscillator, this state would simply rotate in a circle, and retain its uncertainty-limited shape. As can be seen, the quantum quartic oscillator spreads out in phase space, and its state has support over the whole annulus of states with the appropriate energy. Because the Hamiltonian is a polynomial in n , its eigenvalues are integers, and the dynamics are periodic.

As can be seen, the state becomes a cat state with two or three superposed amplitudes at times that evenly divide the period. This is illustrated in Figure 11, where a measure of the area that the state occupies in phase space is plotted over a cycle. The area measure used is $A = \int |\langle \alpha | \psi \rangle|^2 d^2 \alpha / \int |\langle \alpha | \psi \rangle|^4 d^2 \alpha$. Except for the normalisation $\int |\langle \alpha | \psi \rangle|^2 d^2 \alpha = \pi$, this is the same area that was used for the cross sectional of a cloud of trapped atoms in Section 8. These recurrences do not occur classically; their period is an entirely quantum quantity, which is determined by the energy of the ground state.

It is interesting to observe the spiral of zeros that occurs at times when the state is most spread out in phase space, and is furthest from a cat state. There is a set of coherent amplitudes β_i for which $|\beta_i\rangle$ does not overlap the evolved state. When a field of repulsive particles begins in a coherent state, then at later times there are certain amplitudes that it can not have.

Because the quantum and classical dynamics are so different, it is clear that the truncated Wigner method will not work. It might perform reasonably at small times, when the amplitudes are spreading out [43]. However, the classical trajectories do not sample the Wigner distribution once the amplitudes have completed a full cycle. The Wigner distribution of this system becomes negative at the times when the state becomes a cat state, and it ceases to be a probability distribution at these times.

Given the dynamics of this system, it is reasonable to expand it over coherent states, because it occupies a compact area of phase space. The variational amplitudes do not actually have to vary: the state quickly spreads out to cover the whole circle in phase space, so we might as well choose a grid of amplitudes to cover this circle, and vary the coefficients by Dirac-Frenkel or a similar method. In fact, once we have found a stable variational




Figure 10. Phase space components of a coherent state evolving in a quartic potential.

method with coherent states, we might expect the variational amplitudes to spread out as the state does, and then remain fixed as it evolves. The number of coherent states required to cover this area is similar to the number of number states required to expand the oscillator state, and it is much simpler to compute the dynamics in the orthonormal basis of number states.




Figure 11. Area in phase space on which the components are large.

Number states are the easiest way to compute the dynamics of this system in practice, and this is very useful for comparison with the variational results. Despite variation of the amplitudes being fairly pointless when simulating a quartic oscillator, it still makes a good system to test the method. If we start out with amplitudes reasonably close to the initial coherent amplitude, then they need to spread out as the state does, and this is a good system to test whether they do so in a stable fashion.

15. Truncated Wigner dynamics

Sections 7 and 8 described the atomic field exclusively in terms of its order parameter, except for a handful of places where shot noise was mentioned. This description in terms of classical waves ignores the fact that atoms are particles. As discussed in Section 5, it accurately describes a condensate of many bosons whose wave functions are coherent. Chapter 4 will consider atoms in a one-dimensional trap. The atomic field in these traps is nearly coherent, and experiments that have been performed with split traps measure the extent to which the phase of the order parameter is poorly defined. This is the realm of quantum optics, where the atomic nature of matter reasserts itself by causing small fluctuations around the classical dynamics of the order parameter.

Investigating these quantum phenomena requires the computation of the dynamics of a quantum field. This is very hard to do in general. Fortunately, quantum optics is a special case. The fact that the fields are nearly coherent can be exploited to generate an ensemble of classical fields whose motion closely approximates that of the quantum field. There are several ways to do that; we will use the truncated Wigner method, because the ensembles that it generates tend to be small enough to compute with, allowing quantitative predictions

to be made.

To define an ensemble of classical states, consider a function f taking phase space to the real numbers; we will treat it as a probability distribution from which the ensemble is drawn. The Wigner correspondence takes such a function to an operator $\mathcal{W}f$ in Hilbert space—the exact mapping is defined below [44]—and the function f is called the Wigner distribution corresponding to the operator $\mathcal{W}f$. The correspondence is an isomorphism between the L^2 inner product of functions and the Hilbert-Schmidt inner product of operators; a pair of functions f and g satisfies the identity

$$\text{Tr}[(\mathcal{W}f)^\dagger \mathcal{W}g] = \int f^*(\alpha)g(\alpha) d^2\alpha. \quad (29)$$

Wigner distributions exist for a large class of operators [39], and for physical purposes we can consider \mathcal{W} to be invertible. Occasionally, when α is a random variable and the function f is its probability density, we will abuse the notation and write $\mathcal{W}\alpha$ where we strictly mean $\mathcal{W}f$.

The fact that the mapping \mathcal{W} is an isomorphism has a particularly useful consequence. Given a density operator ρ and an observable A , it follows that

$$\langle A \rangle = \text{Tr}[A\rho] = \int (\mathcal{W}^{-1}A)(\mathcal{W}^{-1}\rho). \quad (30)$$

This allows expectation values in the state ρ to be calculated as moments of its Wigner distribution.

There are several equivalent ways to construct the Wigner correspondence. It is defined in terms of power series here; the other ways are described in Section 16, and it will be shown at that point that the correspondence is an isomorphism. We will first introduce the correspondence for the operator e^{a+a^\dagger} , where it takes the simplest form, then generalise it to other operators. This operator expands to a Taylor series of terms in which the operators a and a^\dagger are raised to various powers and ordered in various ways. Collecting terms that comprise the same factors of a and of a^\dagger arranged in different orders, the series becomes

$$e^{a+a^\dagger} = \frac{1}{0!} + \frac{a}{1!} + \frac{a^\dagger}{1!} + \frac{a^2}{2!} + \frac{a^{2\dagger}}{2!} + \frac{a^\dagger a + a a^\dagger}{2!} + \dots + \frac{a^\dagger a^2 + a a^\dagger a + a^2 a^\dagger}{3!} + \dots \quad (31)$$

Each term is the sum of every distinct ordering of a fixed set of factors a and a^\dagger ; any series with these terms, but possibly different coefficients, is called a symmetrically ordered expansion. The function $\mathcal{W}^{-1}e^{a+a^\dagger}$ is constructed from this series by replacing the annihilation operator a with a complex variable α , to give

$$(\mathcal{W}^{-1}e^{a+a^\dagger})(\alpha) = \frac{1}{0!} + \frac{\alpha}{1!} + \frac{\alpha^*}{1!} + \frac{\alpha^2}{2!} + \frac{(\alpha^*)^2}{2!} + \frac{2|\alpha|^2}{2!} + \dots + \frac{3\alpha^* \alpha^2}{3!} + \dots = e^{2\text{Re } \alpha}.$$

In general, the Wigner distribution for any operator is constructed by expanding it in a symmetrically ordered power series, then substituting a complex variable for the annihilation operator and its conjugate for the creation operator. Most operators that occur in calculations can be expanded this way [39].

As a matter of notation, we will define an operation s that puts a product of creation and annihilation operators into symmetric order. For example, $s(a^2) = a^2$, $s(a^\dagger a) = \frac{1}{2}(aa^\dagger + a^\dagger a)$, and $s(a^2 a^\dagger) = \frac{1}{3}(a^2 a^\dagger + aa^\dagger a + a^\dagger a^2)$. The denominator is the number of terms in the symmetric expression. This has the property that $s(a^n (a^\dagger)^m \dots) = W(\alpha^n (\alpha^*)^m \dots)$, where the monomial on the right hand side is obtained by substituting α for a in the product that is symmetrised on the left hand side.

Going the reverse way, the operator $\mathcal{W}f$ is constructed by expanding the function f as a Taylor series, then replacing the term in $\alpha^m (\alpha^*)^n$ by the corresponding term from Equation 31, keeping the same coefficient. This can be extended to systems with several degrees of freedom, in which case the distribution is a function of several complex variables. A set of values for these variables is a point in the multidimensional phase space of the system. When the system is an atomic field, there will be one degree of freedom for each atomic orbital; in simulations, the atoms are restricted to a finite set of orbitals.

It can be shown [40] that the Wigner distribution for a coherent state is

$$(\mathcal{W}^{-1}|\beta\rangle\langle\beta|)(\alpha) = 2e^{-2|\alpha-\beta|^2}.$$

The distribution for the vacuum state is

$$(\mathcal{W}^{-1}|0\rangle\langle 0|)(\alpha) = 2e^{-2|\alpha|^2}. \quad (32)$$

When this is treated as a probability distribution for the complex number α , the stochastic average amplitude is $\overline{|\alpha|^2} = \frac{1}{2}$. This is commonly described as “adding half a virtual atom”, to account for the vacuum energy in the ground state. As an example, the Wigner distributions for the quartic oscillator states of Figure 10 are shown in Figure 12. The oscillations in the Wigner distribution cause the moments that correspond to the position and momentum observables to display the interference fringes from the superposed coherent amplitudes in the state.

Truncated Wigner is a Monte-Carlo method based on this mapping, that simulates quantum dynamics [45, 46, 47]. The Wigner distribution $\mathcal{W}^{-1}\rho_0$ for the initial state ρ_0 serves as a probability distribution, from which an ensemble of classical states $\alpha_1, \dots, \alpha_n$ is drawn at random. This ensemble captures the statistics of quantum observables in the initial state. According to Equation 30,

$$\langle A \rangle = \int (\mathcal{W}^{-1}A)(\mathcal{W}^{-1}\rho_0) \approx \frac{1}{n} \sum_{i=1}^n (\mathcal{W}^{-1}A)(\alpha_i),$$

and the approximation converges as the average is taken over a larger ensemble.

In order to simulate dynamics, each classical state α_i is treated as an initial value for the classical equations of motion, which are solved to give a trajectory of the classical field; the resulting ensemble of trajectories generates an ensemble of states $\alpha_i(t)$ at any later time t , and this is treated as if it had been drawn from the distribution $\mathcal{W}^{-1}\rho(t)$ corresponding to the the state of the quantum field at that time. In particular, it is assumed that




Figure 12. Wigner distribution for a coherent state evolving in a quartic potential.

averages over the ensemble $\alpha_i(t)$ approximate the expectation values of the corresponding symmetrically ordered observables at time t .

In this way, the quantum system is approximated by an ensemble of states of the equivalent classical system, distributed in a way that satisfies the uncertainty principle. This picture is easy to reason about, and it is usually easy to tell whether the ensemble will remain

compact and the averages will converge. If the classical dynamics are linear, the ensemble will not spread out too much, and averages that are converged in the initial ensemble will remain converged as the simulation progresses. The exception would be classically chaotic systems, which could cause the initial bunch of amplitudes to generate a wildly diverging set of trajectories.

The reader is bound to have some doubts about the possibility of carrying out such a procedure, and about how well its predictions approximate those of quantum mechanics. These doubts are justified. Many density operators, including those plotted above for the quartic oscillator, have Wigner distributions that assume negative values in some regions of phase space. Such states are not often used as initial states in this method, because the function $\mathcal{W}^{-1}\rho_0$ is not a probability distribution. However, this could be done by drawing separate ensembles from the positive and negative parts of the function, and weighting the averages by the signs of the samples.

A more fundamental problem is that some Hamiltonians cause states with negative Wigner functions to evolve from a coherent initial state. This is the case for the quartic oscillator. Its truncated Wigner dynamics are shown in Figure 13, and they are very different from the quantum dynamics. The initially Gaussian classical distribution becomes a spiral in phase space, due to the nonlinear frequency increasing with amplitude. The spiral is very thin; it quickly becomes thinner than the pixels of the images, and breaks up into points. The absence of interference fringes make it clear that there are no cat states in these distributions.

The truncated Wigner formalism can be applied to this Hamiltonian, the coherent initial state can be sampled, and averages can be taken over ensembles of trajectories at later times. However, at times when the evolved quantum state has a negative Wigner distribution, these averages can not possibly give the quantum expectation values for all observables. When the dynamics are formally derived in Section 17, an approximation will be introduced that allows this to happen: this approximation is what “truncated” refers to.

The method succeeds to the extent that the motion of the quantum field is approximated by the motion of the classical field. Cold atoms are not at all like classical particles, and applying the truncated Wigner method to their wave function would give unphysical results. On the other hand, a cold atomic field is very like a classical wave, so it is reasonable to apply the method to cold atoms when their state is represented in second quantisation. When this is done, each classical configuration has the same form as the order parameter $\psi(x)$. The initial state of the field is coherent, so the Wigner distribution is a normal distribution, and samples can be drawn by adding Gaussian noise to the initial order parameter according to Equation 32. The classical law of motion is the Gross-Pitaevskii equation; this can be solved by the usual methods, or at least those of them that can handle the irregular noise.

If the system starts in a state where every degree of freedom is occupied by many quanta, and remains in such a state at later times, then there are reasons to expect the ensemble of classical trajectories to accurately sample the Wigner function [12]. Those




Figure 13. Truncated Wigner dynamics for a coherent state evolving in a quartic potential.

reasons will be presented in Section 17.

The truncated Wigner method has been useful in simulating fields in second quantisation, because it is practical to represent the state of a classical field on a computer, but not the full state of a quantum field. This approach has been used for quantum optics [48], and for dilute gas Bose-Einstein condensates [49].

16. Cahill frames for operators

Section 15 introduced the Wigner distribution, a mapping \mathcal{W} that takes phase space functions to Hilbert space operators. This was defined as a correspondence between power series. Functions expand to a power series in α and α^* , and operators to a power series in symmetrically ordered combinations of a and a^\dagger ; the function and the operator correspond when these series have the same coefficients.

There is a different but equivalent way to define this correspondence. Operators over Hilbert space have an inner product $\langle A, B \rangle = \text{Tr}[B^\dagger A]$, known as the Hilbert-Schmidt inner product; for the rest of this section, orthogonality of operators is under that inner product. It is possible to choose an orthonormal basis $T(\alpha)$, with one element for each point in phase space, such that the mapping \mathcal{W} is a basis expansion of the form

$$\mathcal{W}f = \int T(\alpha)f(\alpha) d^2\alpha. \quad (33)$$

The details were worked out in a pair of papers by Cahill and Glauber [39, 40], and this thesis will refer to the operators $T(\alpha)$ as a Cahill basis.

The starting point is the normally and antinormally ordered forms for the displacement operator,

$$D(\alpha) = e^{-|\alpha|^2/2} e^{\alpha a^\dagger} e^{-\alpha^* a} = e^{|\alpha|^2/2} e^{-\alpha^* a} e^{\alpha a^\dagger}. \quad (34)$$

This leads to

$$\text{Tr}[D(\alpha)] = \pi^{-1} \int \langle \xi | D(\alpha) | \xi \rangle d^2\xi = \pi^{-1} e^{-|\alpha|^2/2} \int e^{\alpha\xi^* - \alpha^*\xi} d^2\xi = \pi\delta(\alpha),$$

so the displacement operators are orthogonal, with

$$\text{Tr}[D^\dagger(\beta)D(\alpha)] = \pi\delta(\alpha - \beta).$$

Also, the displacement operators span the space of operators that are bounded under the norm derived from the Hilbert-Schmidt inner product. This includes the density operators, which satisfy $\|\rho\|^2 = \text{Tr}[\rho^2] \leq 1$. The proof is somewhat complicated, and is given in [39]. The result is that the normalised displacement operators $\pi^{-\frac{1}{2}}D(\alpha)$ are an orthonormal basis. The complex Fourier transforms of the displacement operators,

$$T(\alpha) = \pi^{-1} \int e^{\alpha\beta^* - \alpha^*\beta} D(\beta) d^2\beta,$$

also form an orthonormal basis. It can be shown [39] that the Wigner distribution is an expansion over this basis, as in Equation 33.

The Hilbert-Schmidt inner product makes physical sense because it relates states and observables to expectation values, in a similar way as the usual inner product does in Hilbert space. The expectation value of an observable A in the state ρ is $\langle A \rangle = \text{Tr}[\rho A] = \int \mathcal{W}^{-1}\rho \mathcal{W}^{-1}A$. The second equality is an expansion of the inner product $\langle A, \rho \rangle$ over the

basis T . Given an ensemble of amplitudes from the distribution ρ , this moment can be evaluated as the average of $\mathcal{W}^{-1}A$ over those samples.

The Wigner distribution can be generalised by varying the order of a and a^\dagger in the power series and in the frame elements. Equation 34 can be generalised to define a more generally ordered displacement operator,

$$D(\alpha, s) = e^{s|\alpha|^2/2} e^{\alpha a^\dagger - \alpha^* a}$$

Note that $D(1, 0) = \exp(a + a^\dagger)$, the operator whose Taylor series was used in Section 31 to define the terms in a symmetrically ordered power series. It is also the case that $D(1, -1) = \exp(a)\exp(a^\dagger)$ and that $D(1, 1) = \exp(a^\dagger)\exp(a)$, and the expansions of these operators define the terms that occur in normally and antinormally ordered power series. Given an function with a power series

$$P(\alpha) = c_{00} + c_{01}\alpha + c_{10}\alpha^* + c_{02}\alpha^2 + c_{11}\alpha^*\alpha + c_{20}(\alpha^*)^2 + c_{03}\alpha^3 + \dots,$$

a corresponding operator can be defined with the normally ordered power series

$$\mathcal{W}_1 P = c_{00} + c_{01}a + \dots + c_{11}a^\dagger a + \dots + c_{12}a^\dagger a^2 + c_{21}(a^\dagger)^2 a + \dots + c_{22}(a^\dagger)^2 a^2 + \dots,$$

and a similar correspondence \mathcal{W}_{-1} can be defined for antinormal order.

The Fourier transform $T(\alpha, s) = \int e^{\alpha\xi^* - \alpha^*\xi} D(\xi, s) \pi^{-1} d^2\xi$ can be defined for arbitrary values of s . In particular, $T(\alpha, -1) = |\alpha\rangle\langle\alpha|$, so the expansion $P(\alpha) = \text{Tr}[T(\alpha, 1)\rho]$ is the Glauber-Sudarshan distribution, and $Q(\alpha) = \text{Tr}[T(\alpha, -1)\rho]$ is the Hushimi distribution. The set $T(\alpha, s)$ fails to be a frame in the same way that $D(\alpha, s)$ does. For example, the Hushimi distribution for a position state, $Q(\alpha) = |\langle\alpha|x\rangle|^2$, forms a stripe in phase space parallel to the momentum axis, and $\|Q\|$ is infinite. However, the expansion of an operator A over the $T(\alpha, s)$ has a reasonable norm provided that the length and wave number scales involved in the operator are close to the characteristic ones for the harmonic oscillator whose annihilation operator is a .

The set of $D(\alpha, s)$ is dual to the set $D(\alpha, -s)$ in the sense that $\text{Tr}[D^\dagger(\beta, s)D(\alpha, -s)] = \pi\delta(\alpha - \beta)$. There is a similar dual property $\text{Tr}[T^\dagger(\beta, s)T(\alpha, -s)] = \pi\delta(\alpha - \beta)$ for the Fourier transformed operators, which allows expectation values to be computed by sampling. This involves sampling the distribution $\mathcal{W}_s\rho$ in phase space, then evaluating the values at the sample points of the function $\mathcal{W}_{-s}A$.

These expansions are quite neat from a formal point of view. In order for them to be useful for calculations, the power series must all converge, and the distribution functions must be well behaved. It turns out [39] that almost every operator is the sum of a normally ordered power series that converges, but very few are the sum of an antinormally ordered power series.

This forces a compromise in the use of the Monte Carlo methods. In order to evaluate moments, we need to choose a value of s such that the observable A corresponds to a well-behaved function $\mathcal{W}_{-s}A$, and also so that the density operator has a distribution $\mathcal{W}_s\rho$ that

we can sample. In practice, almost all simulations start with coherent states, which correspond to normal distributions in all the representations. So the choice of s is determined by the requirement that the observables correspond to useful functions. Of course, a method that can only represent a limited set of density operators must compromise on the accuracy with which it represents dynamics, because the simulated density operator must remain in the representable set. We will return to this in Section 19 below.

Historically, much of the early work in quantum optics used the P distribution \mathcal{W}_{-1} . This had the advantage that the observables needed to be expanded as normally ordered power series, which made intensity and photon shot noise quite simple to calculate. However, it was quite lucky that the system states could be represented in this way. The reason it worked is that the first systems being treated were lasers, or occasionally stars, and the states of light being considered were thermal or coherent ones [42]. These are about the only states whose P distributions are ordinary functions that can be sampled for Monte-Carlo methods. When quantum optics was applied to nonlinear systems such as parametric oscillators, other methods had to be employed. These are the systems that exhibit the really interesting quantum optical effects such as violations of Bell inequalities.

17. Phase space dynamics

In quantum optics, a field where phase space simulations are often used, the systems being studied are usually open ones, whose law of motion is a quantum master equation. Although the systems in this thesis are closed, and their states satisfy Schrödinger's equation, this section describes the full master equation for completeness.

The Hamiltonian motion of a density operator $\rho(t)$ is given by

$$i\hbar\rho'(t) = [H, \rho].$$

The master equation extends this with various terms that account for particles being exchanged with an environment [50], with a result similar to

$$i\hbar\rho'(t) = [H, \rho] + \frac{\gamma}{2}(N+1)(2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) + \frac{\gamma}{2}N(2a^\dagger \rho a - a a^\dagger \rho - \rho a a^\dagger).$$

Here γ is the rate at which the system equilibrates with the environment, and N is the number of particles that would occupy a part the environment that is similar to the system. Because the mappings \mathcal{W} and \mathcal{W}_s are linear, the terms can be treated in isolation, and the results added back together as an equation of motion for $\mathcal{W}^{-1}\rho$. In order for this equation to be solved efficiently by a Monte-Carlo method, the derivatives should be at most second order, and their coefficients need to form a positive definite matrix [51].

The second and third terms are not problematic. For example [42],

$$\mathcal{W}^{-1}(2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) = \left(\frac{\partial}{\partial\alpha}\alpha + \frac{\partial}{\partial\alpha^*}\alpha + 2\frac{\partial^2}{\partial\alpha\partial\alpha^*} \right) \mathcal{W}^{-1}\rho$$

The derivation of the differential operators corresponding to irreversible master equation terms is a fairly tedious exercise in integration by parts; however, the irreversible terms are

restricted to a handful of forms, and the derivations can be found in textbooks. The ability to treat these terms has been the main reason for interest in phase space methods.

Hamiltonian dynamics are less straightforward. In classical statistical mechanics, the equivalent of the density operator is simply a probability distribution $P(p, q)$ over phase space, whose Hamiltonian evolution is given by the Liouville equation

$$\partial_t P(t, p, q) = \{H_c, P\} = \partial_q H_c \partial_p P - \partial_p H_c \partial_q P. \quad (35)$$

The bracket here is a Poisson bracket, and the classical Hamiltonian $H_c = \mathcal{W}^{-1}H$ is a function of the canonical coordinates. Equation 35 can be solved by a very simple Monte-Carlo method. An ensemble of states (q_i, p_i) , drawn from the distribution $P(p, q, 0)$, will continue to be distributed according to $P(p, q, t)$ provided that each state follows the classical equations of motion $\dot{p} = \{H_c, p\}$ and $\dot{q} = \{H_c, q\}$.

For unitary quantum dynamics, an equation similar to Equation 35 can be derived [24] for the Wigner distribution $W(t, p, q) = \mathcal{W}^{-1}\rho(t)$. This is

$$\begin{aligned} \partial_t W(t, p, q) &= \{H_c, W\}_M \\ &= \frac{2}{\hbar} \sin\left(\frac{\hbar}{2}(\partial_p W \partial_q H - \partial_p H \partial_q W)\right) H_c(p, q) W(t, p, q) \\ &= \{H_c, W\} + O(\hbar). \end{aligned}$$

The bracket here is called a Moyal bracket; the partial derivative $\partial_{.W}$ acts only on the Wigner distribution, and $\partial_{.H}$ acts only on H_c . This reverts to a Poisson bracket in the high energy limit where the effects of the derivative operators are large compared to \hbar . In general, it includes a series of third, fifth, and all odd order derivatives. In the special case of a harmonic oscillator Hamiltonian, the third and higher order derivatives of H vanish, and this reverts to the classical equation of motion; a Monte-Carlo Wigner method can solve the harmonic oscillator exactly. When the potential is not harmonic, it is necessary to discard these higher order derivatives. This is why the method is called truncated Wigner.

The Moyal bracket gives the equation of motion for the Wigner distribution. The other distribution functions in the Cahill family are linearly related to the density operator; for a fixed Hamiltonian H , there must be some differential operator whose effect on the distribution function is equivalent to the effect of commutation with H on the density operator. However, it is not clear that this differential operator can be derived from the classical Hamiltonian H_c , or a phase space function corresponding to H in any order. As far as I know, it is an open question whether generalised Moyal brackets can be found such that $\{H_s, W_s\}$ generates an equation of motion for the s -ordered distribution W_s that is equivalent to the quantum Liouville equation.

In second quantisation, the degrees of freedom are single-particle orbitals, and the quantum numbers count the atoms. The truncation of Moyal brackets to Poisson brackets is valid in the limit of many atoms per orbital. In order to increase this, it is desirable to reduce the system of trapped atoms to as few orbitals as possible before introducing classical configurations. In Section 32 below, we will reduce the split condensate to a set of one

dimensional order parameters $\psi_i(z)$ for different orbitals of the transverse trap. Presumably this is as far as the orbitals can be reduced, and it would be the ideal model on which to base a Wigner simulation. However, these one dimensional modes are complicated to implement, and the results in this thesis are from a two dimensional simulation. Section 33 will argue that the unoccupied modes introduced by this are isolated from the occupied ones, and that they should not affect the results very much.

18. Positive P representations

The positive P distribution [52] compromises in another direction. Instead of expanding operators over phase space, with components of a set of Cahill operators, it expands them over a double phase space, with components of the operators

$$T_+(\alpha, \beta) = \frac{|\alpha\rangle\langle\beta^*|}{\langle\beta^*|\alpha\rangle}.$$

These are a superset of the operators in the Glauber-Sudarshan P distribution,

$$T_{-1}(\alpha) = |\alpha\rangle\langle\alpha|.$$

When these operators are expanded over position states by Equation 5, they become two-dimensional Gaussian wave packets,

$$\langle x|T_+(\alpha, \beta)|x'\rangle = \frac{1}{\sqrt{\pi}l_0\langle\beta^*|\alpha\rangle} \times \exp\left(-\frac{x}{l_0} - \sqrt{2}\operatorname{Re}\alpha\right)^2 - \frac{x'}{l_0} - \sqrt{2}\operatorname{Re}\beta\right)^2 + i\sqrt{2}\operatorname{Im}\alpha x/l_0 + i\sqrt{2}\operatorname{Im}\beta x'/l_0), \quad (36)$$

so they form a frame for operators, very similar to the multimode Gaussian wave packets that will be analysed in Chapter 5.

Because the operators $T_+(\alpha, \beta)$ form a frame, expansions over them are not unique. This freedom turns out to have a very useful consequence: every density operator has an expansion with positive coefficients,

$$\rho = \mathcal{W}_+ P_+ = \int P_+(\alpha, \beta) T_+(\alpha, \beta) d\alpha d\beta,$$

where the values of P_+ are positive and real. This expansion is typically not unique, but one form of it can be derived by inverting the Husimi distribution, as

$$\rho = \int T(\gamma) Q(\gamma) d^2\gamma, \quad (37)$$

where

$$T(\gamma) = \frac{1}{\pi^2} \int \frac{|\gamma + \delta\rangle\langle\gamma - \delta|}{\langle\gamma - \delta|\gamma + \delta\rangle} e^{-|\delta|^2} d^2\delta = \frac{1}{\pi^2} \int |\gamma + \delta\rangle\langle\gamma - \delta| e^{2i\operatorname{Im}\delta^*\gamma} d^2\delta.$$

An expansion of the same form Equation 37 holds with $T(\gamma, 1)$ in place of $T(\gamma)$, although it is not immediately clear whether these operators are the same.

Because T_+ has the same form as T_{-1} , it satisfies a similar inner product relation as does T_1 , and

$$\text{Tr}[AW_+P_+] = \int \mathcal{W}_-A(\alpha, \beta)P_+(\alpha, \beta) d\alpha d\beta,$$

where the function \mathcal{W}_-A is obtained from the normally ordered power series for the operator A by replacing a with α and a^\dagger with β^* .

Furthermore, it was shown in [52] that the commutator $[H, \rho]$, at least when the Hamiltonian H is quartic or lower order in a and a^\dagger , corresponds to the application of a second-order differential operator to the distribution P_+ , and that the diffusive part of this operator has a positive definite matrix. This is derived by equating expressions such as $aT_+(\alpha, \beta)$ and $T_+(\alpha, \beta)a^\dagger$ with the effects of certain differential operators on T_+ . If such derivations can be carried out for the positive P distribution by any method as general and clear as the Moyal bracket, that method is yet to be discovered.

This allows the dynamics of such Hamiltonians to be computed as a Monte-Carlo ensemble of trajectories, in the method of Section 19. The trajectories are not classical configurations, but pairs of configurations, corresponding to the phase space amplitudes α and β . Because these dynamics satisfy a Fokker-Planck equation, the solution for P_+ is guaranteed to remain a normalised distribution. The operators T_+ are a frame, so in general the coefficients of expansions could become arbitrarily large.

Perhaps the most impressive demonstration of this method was a simulation of the quartic oscillator [43]. The system simulated was the quartic oscillator, described in Section 14. Its dynamics generate Schrödinger cat states, whose Wigner distributions assume negative values, so there is no chance of calculating the dynamics by a truncated Wigner method. Starting from a coherent state, the classical trajectories spread from a disk into a spiral as the states with higher amplitudes rotate at higher frequencies, and the distribution ends up very similar to a Poisson distributed mixture of number states. But when this system is simulated by positive P, the expectation values of the oscillator quadratures reproduce those of a quantum quartic oscillator, indicating that the cat states turn up when they are supposed to.

In principle, this method is perfect for simulations. The downside is its convergence. A distribution function that comprises the coefficients of a state along an orthonormal basis is unique. Therefore, the methods based on Cahill bases guarantee that, when the simulated quantum state has support in a limited region of phase space, the Monte-Carlo trajectories will remain in this region: they are distributed according to the unique expansion. The frame of Gaussian wave packets allows a range of expansions for any operator, and it usually turns out that many of these are positive distribution functions. Some of these expansions over pairs of configurations can spread out much farther in phase space than does the expansion of the same state over a Cahill basis; this means that many more samples must be drawn from the distribution before the averages converge to expectation values. And it is usual for the positive P distribution to spread out in this way over the course of a Monte Carlo simulation.

The form of Equation 36, in which a positive P distribution is derived from the Hushimi distribution, hints at these convergence problems. Consider the pure states $|\alpha\rangle \pm |\beta\rangle$ and the mixed state $|\alpha\rangle\langle\alpha| + |\beta\rangle\langle\beta|$ and ignore normalisation. The Hushimi distribution for the mixed state is $Q_0(\gamma) = \exp|\gamma - \alpha|^2 + \exp|\gamma - \beta|^2$. For the pure states, it has an extra term

$$Q_{\pm}(\gamma) = Q_0(\gamma) \pm 2e^{-\frac{1}{2}|\gamma - \alpha|^2 - \frac{1}{2}|\gamma - \beta|^2} \cos \text{Im}(\alpha^* - \beta^*)\gamma.$$

The size of the extra term is bounded by $2e^{-\frac{1}{2}|\alpha - \beta|^2}$, and the ripples that it adds to the distribution have wave number $|\alpha - \beta|$. In principle, the quantum state can be recovered from the Hushimi distribution; however, in some states, an enormous number of samples will be required before one could discern the ripples in their distribution and know whether the state was a coherent superposition or a statistical mixture.


This is illustrated in Figure 14, for $\alpha = 1.65 = \pm\beta$. The amplitudes $\langle\gamma|(|\alpha\rangle \pm |\beta\rangle)$ are plotted on the left, and their squared moduli, the Hushimi distribution, on the right. The values of α and β have been chosen to show the small ripples in the Hushimi distribution that denote the phase of the superposition. This phase of the superposition can be determined in principle from the Hushimi distribution, but it is much more clearly discernable from the phase of the complex amplitudes; and it remains just as clear when α and β become large, and the ripples become exponentially shallower. In the limit where the states $|\alpha\rangle$ and $|\beta\rangle$ become orthogonal, their Hushimi distributions become identical.

If a Wigner distribution is used instead of a Hushimi distribution, the phase of the superposition is still encoded in fringes in the same region of phase space. However, for a Wigner distribution, these fringes remain large as the difference $|\alpha - \beta|$ increases. This is a significant advantage of the mapping from operators to Wigner functions being an isomorphism: when $|\alpha - \beta|$ is large, the even and odd superpositions are nearly orthogonal, so their Wigner distributions must be nearly orthogonal under the L^2 inner product. It is not possible for very different states to have nearly identical Wigner distributions.

For some systems, the positive P distribution spreads out so rapidly that the stochastic dynamics are no longer equivalent to the quantum dynamics [53]. The equivalence between the Fokker-Planck equation and the quantum law of motion relies on the distribution function decaying rapidly at large amplitudes, so that differential operators can be moved from the T_+ operators to the distribution function using integration by parts. It is always possible to derive a Fokker-Planck equation by assuming that the distribution decays rapidly enough, however there are some nonlinear systems in which the solutions to this equation spread out too far for the assumption to hold. These systems usually have stochastic equations of motion in which some of the sample amplitudes will diverge to infinite amplitudes at finite times.

Even when the boundary terms remain zero, so that the phase space dynamics are formally equivalent to the quantum dynamics, it is often the case that very many trajectories must be sampled over a large region of phase space in order for the averages to converge within the small region of phase space where the state is located. Typically, the distribution remains dense at short times, but the trajectories spread out as they propagate, and become

Symmetric superposition $|\alpha\rangle + |\beta\rangle$



Asymmetric superposition $|\alpha\rangle - |\beta\rangle$




Figure 14. Modulus and phase of phase space components of cat states.

sparsely distributed. Many systems of interest require an impractically large number of samples to simulate accurately.

There are various ways to improve the convergence. Piotr Deuar used the extra freedom in the frame expansion of ρ to change the Fokker-Planck equation, so that it had a smaller diffusion term, and the trajectories in some systems stayed converged for longer. This allowed some simulations of a quartic oscillator, which any classical dynamics are bound to get wrong. The positive P distribution has been extended to fermions, using Gaussian operators instead of coherent state projectors [54], and a Japanese group developed this method to simulate Fermi-Hubbard systems [55]. These were the largest Fermi-Hubbard systems for which quantitative results are known.

19. Monte Carlo dynamics

The phase space methods convert the quantum Liouville equation into a partial differential equation for a probability distribution over phase space; this could be the P function, the Wigner distribution, or the positive P distribution. For a quantum field, the phase space has a large enough dimension that solving this PDE on a grid is out of the question. However,

it is sometimes possible to find an ordinary differential equation $\alpha'(t) = f(\alpha(t))$ with the property that an ensemble of trajectories $\alpha_i(t)$, each of which satisfies this ODE, has a distribution $P(t, \alpha)$ that satisfies the PDE. This allows a numerical solution to be found efficiently. An set of initial values $\alpha_i(0)$ is drawn from the distribution $P(0, \alpha)$, then each trajectory $\alpha_i(t)$ is computed independently by solving the ODE. Because these computations are independent, they can be carried out very efficiently on a parallel computer or a cluster.

However, only partial differential equations with a special form can be unravelled in this way. If the ordinary differential equation is derived from a Hamiltonian, then the distribution function for an ensemble of trajectories satisfies the Liouville equation. This is a special form of partial differential equation; for example, it only has first order derivatives in α .

A broader class of partial differential equations can be unravelled if the ordinary differential equation is generalised to a stochastic differential equation of the form $\alpha'(t) = f(\alpha(t)) + \xi(t)$, where ξ is a white noise function, the derivative of a random walk. When an ensemble of trajectories satisfy an equation of this form, their distribution satisfies a second-order partial differential equation of the Fokker-Planck form. The irreversible terms mentioned in Section 17 result in a second order PDE, so the dynamics of an open system almost always unravels to an ensemble governed by a stochastic differential equation instead of an ordinary one.

The fields that will be treated in Chapter 4 have Hamiltonian dynamics, so there will be no second order derivatives in the equation of motion for the Wigner distribution. However, the Moyal brackets include partial derivatives of third and higher orders. The possibility has been investigated of unravelling these in a phase space of double dimension, in a similar way to the positive-P method [56]. This turned out to be more complicated than the positive-P method, and to suffer from similar convergence issues. It was not pursued very far, and in practice these derivatives are always discarded; when they are, the equation of motion for the phase space distribution reduces to the Liouville equation, which describes the distribution of an classical ensemble, and can be unravelled as a set of independent trajectories each of which satisfies the classical equation of motion.

In Chapter 5, we will consider ensembles of trajectories whose dynamics are not independent. In principle, these can avoid the need to truncate higher order derivatives, and generate dynamics that are closer to the quantum law of motion. However, the interactions between the trajectories can cause calculations of their motion to become numerically unstable, so the calculations are more complicated than those for independent trajectories.

20. Chemistry methods

Some simulation methods, that are either identical to or at least very similar to the phase space ones used in quantum optics, were invented independently by quantum chemists in the 1970s. These were restricted to Hamiltonian dynamics, which suffices for chemical reactions, events which happen too quickly for the environment of the molecules to have much effect while the quantum motion is in progress. The main proponent of these methods

was Heller, who independently derived the truncated Wigner method [45], and analysed in some detail the ways it can fail when quantum superpositions and interference become important.

Heller developed a family of methods that used coherent and squeezed states in different ways. The chemists called these states Gaussian wave packets. One approach is very similar to the truncated Wigner method [13]. This involved expanding the initial quantum state as a superposition of coherent states; such expansions are analysed in detail in Chapter 5 of this thesis. Each of these states was propagated independently, by the following method, and the resulting states at the final time were superposed.

The Hamiltonian was expressed in terms of the position and momentum operators x and p . It was assumed to be quadratic in p . The potential energy was expanded as a Taylor series in x about the centre of the wave packet, and this series was truncated to second order. The result was an effective Hamiltonian that depended on the position of the wave packet. As the wave packet moved, this became a time-dependent Hamiltonian. However, the effective Hamiltonian was quadratic, so a coherent initial state remained a coherent or squeezed state at future times. Later, this method was modified to require the states to be coherent instead of squeezed [17].

Truncating a Taylor series for the potential energy at second order has the same effect as discarding the higher derivatives higher than second order in the Moyal bracket. The Moyal bracket with a quadratic Hamiltonian is equivalent to the Poisson bracket. That is what Heller did in the first variational Gaussian simulations. So the only difference between this and the truncated Wigner method is that it is based on kets instead of density operators.

Heller also looked at a somewhat different approach, in which the entire superposition of coherent or squeezed states is propagated by a variational method [14, 57]. This approach is discussed at length in Chapter 5.

21. Summary

Except for the variational methods introduced by Heller, every phase-space method is derived from a mapping between quantum states and functions over classical phase space. When the methods are used to simulate quantum fields, in second quantisation, the phase space has one complex dimension for every mode over which the field amplitude is expanded. States of fields are mapped to functions of several complex variables, one variable for each mode.

The quantum states in these mappings are density operators, not wave functions. This means that linear combinations of phase space functions correspond to mixtures of quantum states, while superpositions of states correspond to very nonlinear combinations of distributions; the cat states discussed in Section 18 are an example of this. Phase space methods arose from quantum optics, to understand noise in lasers. Lasers are open systems, because the laser cavity has to be coupled to the beam it produces. Therefore their quantum states are mixed, and the phase space methods were designed with that in mind.

Optical fields typically have many modes, and their phase spaces have many dimen-

sions. In this context, the only practical way to discretise the phase space function is to sample phase space amplitudes at random. The methods are designed so that the phase space functions are probability distributions, or at least can be approximated by probability distributions. This allows the functions, and therefore the quantum states, to be discretised as an ensemble of phase space amplitudes that has been drawn from that distribution.

There are three requirements that a phase space distribution must satisfy, in order for it to be useful in simulations. Firstly, the phase space functions corresponding to the states that occur in the system must be probability distributions. Secondly, the dynamics of the distribution must be of a form that allow it to be unravelled into an ensemble of stochastic amplitudes, or vectors of amplitudes for a multimode system, that evolve independently. In practice this means that that distribution must satisfy a Fokker-Planck equation. Finally, the observables of the system that are of interest must correspond to moments of the distribution that converge when a reasonably small number of trajectories are averaged.

For simple systems, with Hamiltonians that are quadratic in the field amplitudes, and some dispersion, all these requirements can be satisfied. These conditions are satisfied by the laser, which was the inspiration of the methods, and the first system that they succeeded in describing.

There are some obvious ways that the classical dynamics of the single phase space methods have to fail. In the truncated Wigner method, the ensemble of trajectories follows classical dynamics. There are some systems whose classical dynamics is entirely different from their quantum dynamics, for example nucleons tunnelling out of the nuclear well potential during radioactive decay.

Suppose that a decaying nucleus, with a high tunnel barrier, was treated by a truncated Wigner method, in which the multiparticle wave function for the nucleons was converted to its Wigner distribution, and the classical dynamics of the nucleons was computed. The energy of the nucleons would be conserved in each classical trajectory, although different trajectories might have more or less energy depending on their random initial conditions. A few trajectories would have enough energy to cross the barrier, and these would escape almost immediately. The rest would remain bound for ever, and in the long time limit the nucleus would have a high probability of remaining intact. This is different from the quantum dynamics, where the nucleus is certain to decay eventually.

In principle, the double phase space methods would get around this. I am not aware of anyone attempting a positive P simulation for a system that tunnels. It is possible that the convergence would fail right at the tunnelling time.

The final chapter of this thesis will consider another way to simulate dynamics in phase space. Instead of looking for a set of trajectories where the phase space amplitudes evolve independently, they will be allowed to interact with each other, and the amplitudes will be treated as variational parameters. The goal will be to find equations of motion for these parameters that most closely approximate the Hamiltonian motion of the state. This ensures that the dynamics actually approximate quantum dynamics. However, this turns out to have serious limitations for simulations of fields. Expanding over coherent states in phase space

is not much better than expanding wave functions over a grid: if the phase space has many dimensions, it is likely that the density operator ρ will overlap with the coherent states in a large volume of phase space, and an unreasonable number of trajectories will be required to cover that volume.

Section 5 discussed how bosonic atoms condense at low temperature. In the absence of interactions between the atoms, there is a phase transition; below a critical temperature, a macroscopic fraction of the gas occupies the ground orbital. This gas becomes a fully coherent atomic field, whose amplitude is an order parameter for the phase transition. When the atoms interact, their positions become correlated, which reduces the interaction energy. But the state remains almost coherent when the interactions are small.

Such condensed atoms are described quite accurately by a classical field theory for the order parameter. Their dynamics can be interesting [58, 59], and thermal and quantum effects appear in the fluctuations [60]. But for the most part these systems are well understood.

Instead of containing the atoms in a volume of space, it is possible to confine them to a line [61, 62]. In this situation, where the atoms are free to move along the line but are confined to it so tightly that they lack the energy to occupy transverse excited states, the atomic field is essentially one dimensional, and its dynamics are more complicated and less well understood [7]. As discussed in Section 5, a condensate does not form in this situation. There are too many orbitals available for a macroscopic fraction of the atoms to occupy the same orbital at any finite temperature.

There are two distinct phases for a cold gas of bosons in one dimension. These depend on the Lieb-Liniger parameter that was introduced in Section 8, the ratio of the healing length to the atomic separation. When the healing length is shorter than the distance between adjacent atoms, there is enough repulsion energy available to confine each atom to a section of the axis that does not overlap those of the adjacent atoms. This phase is called a Tonks gas. The wave function of each atom is energetically prohibited from overlapping the others, so the atoms are constrained to obey an exclusion principle, and this phase is similar to a degenerate gas of fermions.

The other phase, where there are many atoms per healing length, is called the quasi-condensate regime [63]. In this case, density fluctuations are suppressed by the repulsion, but the phase of the atomic field is still free to fluctuate. This gas is similar to a condensate, but it has a finite coherence length.

Several experiments [28, 64] have investigated the coherence of Bose gasses confined in one dimension. Coherence is demonstrated by interference, and to see interference requires the gas to be split between two orbitals. The result is very like a double slit ex-

periment. The atoms line up along interference fringes, and the phase of these fringes is determined by the phase difference between the order parameters in the two clouds. If the phase difference varies along the length of the clouds, the fringes will meander.

As well as phase differences between the two lines, there are also number fluctuations. These are suppressed in a line of condensate, but when the line is split in two, shot noise due to atoms being discrete particles causes some difference in density between the two lines. At later times, the dynamics of the condensate mix this density noise into the phases of the two condensate components, and the coherence length of their relative phase decreases. This causes the straight interference fringes that result from splitting a line of condensate to become meandering as they are held in the trap between splitting and release.

Some of these experiments can measure fluctuations in atomic number directly [64], and compare these to the phase fluctuations.

22. Split trap

Atomic double slit experiments have been carried out with several different types of apparatus. The first experiments used an atom chip, a mirror inside a vacuum chamber with electric wires microfabricated on its surface. Atoms near the surface can be lit by directly incident laser beams from their exposed side, and back lit by the reflections of the beams in the mirror; this suffices to cool the atoms. Near the chip surface, the wires generate a magnetic field with a large gradient, which can form a tight magnetic trap. A straight wire generates a cylindrical trap parallel to itself, and this can be tight enough to confine the atoms to their transverse ground state [65].

The experiment on an atom chip used a gas of ^{87}Rb , whose parameters and geometry [28] are shown in Figure 15. The transverse lengths have been exaggerated for clarity; the actual cloud was longer and thinner than the shape in the figure. The gas was initially trapped along a single line, about $100\ \mu\text{m}$ away from the chip, then it was split into a pair of parallel lines. These lines sat in a transverse double well, the potential that was generated when microwaves radiating from an antenna wire on the chip repelled the gas from the trap centreline while the original trap confined it [66].

In order to investigate the spatial coherence of the of gas, the trap was switched off, allowing the parallel lines of gas to expand and overlap. The resulting interference fringes were photographed, as shown in Figure 16.

The initial trap was very nearly harmonic; its transverse frequency was $2.1 \times 10^3\ \text{Hz}$, and $11\ \text{Hz}$ along the axis, corresponding to characteristic lengths of $0.24\ \mu\text{m}$ and $3.3\ \mu\text{m}$. The transverse shape of the order parameter is very nearly the ground orbital of the trap, while the longitudinal shape has a Thomas-Fermi form and is about 60 times longer than the ground orbital. The thermal energy at the experimental temperature of $78\ \text{nK}$ is somewhat less than the transverse vacuum energy. The temperature can not be set much lower, because the goal is to investigate spatial coherence, and the thermal coherence length needs to be short compared to the trap.

As shown on the diagram, the scales of the scattering length, the atomic separation

All lengths in micrometres Vertical scale differs from horizontal

Atomic mass 86.91, dispersion time scale $1 \mu\text{m}^2 = 1.368 \text{ ms}$

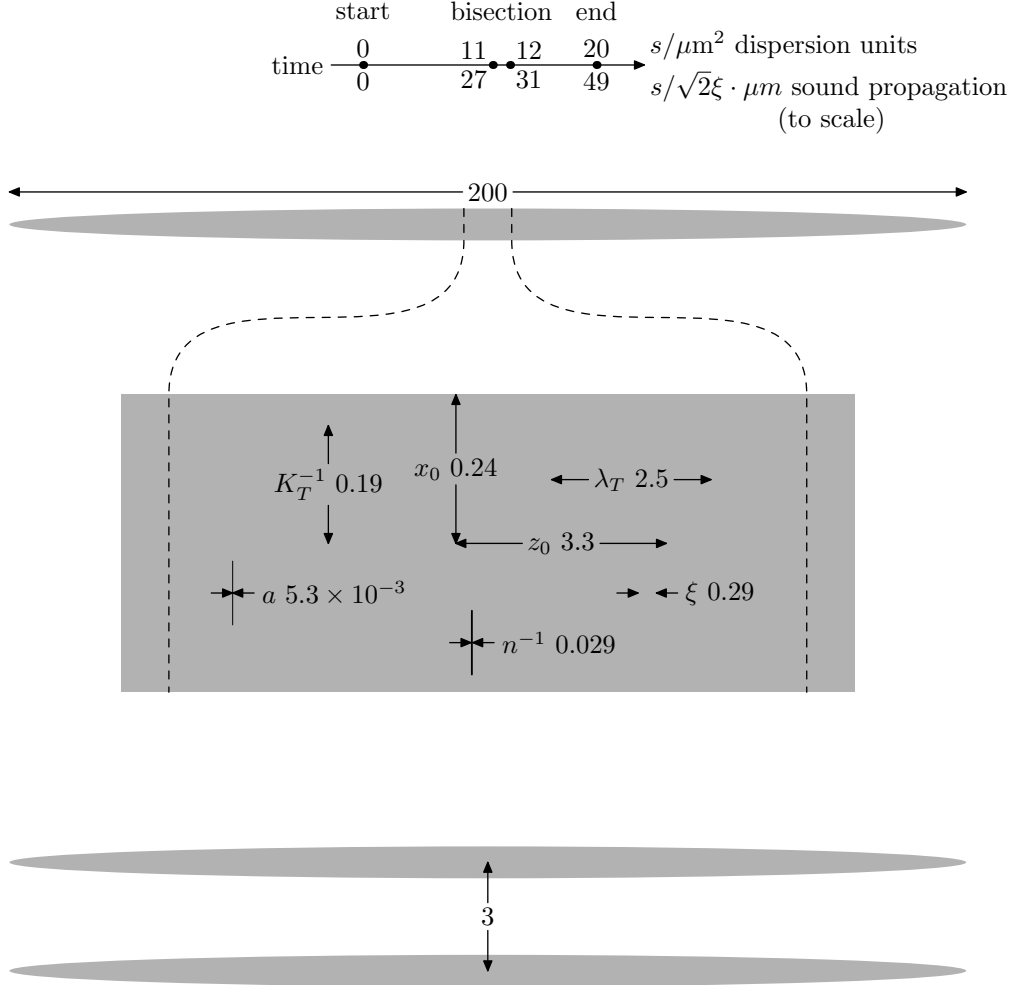


Figure 15. Length, time, and energy scales for the Vienna split trap.

along the axis, and the healing length form a hierarchy, with around an order of magnitude increase between each level. The Lieb-Liniger parameter is around $\gamma = 5 \times 10^{-3}$, so this is well inside the quasicondensate regime.

The timeline shows the start of the trap bisection, marked “start”, the time when the condensate splits into two components, the time the trap potential arrives at its final double well form, marked “bisection”, and the time after the trap split that the coherence length became steady, marked “end”. The time axis is labelled with the dispersion units from Section 3, but it is scaled by the Bogoliubov speed of sound $1/\sqrt{2}\xi$, which will be introduced below in Section 26. Distance along the time axis indicates how far along the trap axis a small excitation of the condensate could propagate in that time.




Figure 16. Interference fringes from two coherent atomic clouds.

The trap was narrow enough that the atoms only moved $1.5 \mu\text{m}$ along the splitting axis. Over the splitting time, this gives a scaled velocity $1.5 \mu\text{m}/12 \mu\text{m}^2 = 0.13 \mu\text{m}^{-1}$, which is adiabatic compared to the transverse wave number $1/x_0 = 4 \mu\text{m}^{-1}$.

The experiment investigated the spatial coherence of the atomic field. This can be determined from the shape of the interference fringes; the phase difference between the lines determines the offset of the fringes, which meander when their phases vary along their length. The phases are limited by shot noise, and the fringes meander in a different way every time the experiment is repeated. Once enough sets of fringes have been photographed, the coherence of the fields can be determined from the correlations between the phases of the fringes at different positions along the line.

The way to observe phase differences is through interference; a double slit experiment compares the phases of the fields at adjacent points on both slits. The phase of a one dimensional condensate fluctuates along the line of atoms, but these experiments never directly compare different points along the trap axis: when the atoms are released, they rapidly expand in the plane transverse to the axis, but they can not move very far along the axis before they are photographed. Phase fluctuations that are common to both components after the trap is split are not discernable from the interference fringes. Only the relative phase of the two lines is visible.

However, the experiments did compare the split gas to two clouds whose relative phase is totally incoherent, with the result shown in Figure 17. The microwave beam can be switched on while the atoms are cooling, a configuration that produces two separate clouds along the parallel lines of the split trap. The phases of these components are independent; all of the phase fluctuations in both lines contribute to the phase difference between the lines. When this was done, the fringes meandered more than when a single cloud was split, and the fluctuations in the common phase could be inferred. As described in Equation 13, this coherence length is

$$\lambda_\tau = 2n/K_T^2 = 2.5 \mu\text{m}.$$




Figure 17. Interference fringes from two incoherent atomic clouds.

The most interesting part of the experiment came when two lines of gas were formed by splitting one cloud, but these were held in the split trap instead of being released immediately. For short holding times, this produced interference fringes which were straight lines, but they meandered more as the gas was held longer, as shown in Figure 18. This increase plateaued after 10 ms, as shown on the time line. The further effect of holding the gas for up to 200 ms was too small to be discernable. But the fringes remained much straighter than those produced by lines of gas that cooled in isolation: the scale on which they meandered was an order of magnitude greater.




Figure 18. Delayed interference fringes from two coherent atomic clouds.

These phenomena have been described using a Luttinger liquid model. The Luttinger liquid is a model of a one-dimensional field with a linear dispersion relation. In a one-dimensional Bose gas with two identical components, the difference between the quantum fluctuations of the two components is isolated from other degrees of freedom, and its

dynamics are described by a Luttinger liquid [8]. Once the trap has split, it comprises two spinless one-dimensional fields. However, because of the tight transverse trap and the repulsion of the atoms, the density in the two components will be nearly equal. If it is assumed that these densities are exactly equal, the only degree of freedom is the relative phase of the components, which can be identified with a single spin- $\frac{1}{2}$ field. This model predicts a final effective temperature such that, in scaled units,

$$K_T^2 = 2\xi^{-2},$$

and the thermal wavelength is close to the one-dimensional healing length. The experimental coherence lengths were close to this, except that the inferred temperature was consistently about 10% higher than the prediction [28].

There is a simple explanation of why the phase fluctuations that result from shot noise in the trap splitting are identical to those that would occur in thermal equilibrium at some effective temperature [28]. The fluctuations occur in quasiparticle modes with long wavelengths, whose excitation energy is small compared to the effective temperature. Given this condition, thermal equilibrium results in equipartition of energy, where each degree of freedom has thermal energy of about kT . For entirely different reasons, the shot noise that results from splitting the trap is white noise, and the intensity of the fluctuations is independent of wavelength. The dispersion relation is linear for long wavelength modes in a cold Bose gas [67, 5], so this constant distribution over wavelength is also a constant distribution over energy, identical to that from a thermal state.

This model has an obvious limitation, in that it only makes sense after the trap has split. It does not allow the splitting process to be analysed in detail, and requires some assumptions to be made about the state of the atoms immediately after the trap splits. These assumptions are the most likely cause of the temperature discrepancy. This chapter investigates ways to analyse the dynamics of the field as the trap splits. This is discussed in Section 24, along with the main limitations of the model.

These results suggest that however long the gas is held, the phases of the two components remain partly correlated; the system never reaches its thermal equilibrium state, in which the phases of the two lines are independent. Instead, the atoms partially equilibrate, then remain in a steady state. When the atoms are released from this steady state, they form interference fringes with the shape expected for atoms released from a thermal equilibrium state, but for equilibrium at an unphysical temperature. The temperature inferred from these interference fringes is an order of magnitude lower than the atoms could have plausibly have.

The maintainers of the trap have provided numerical samples of the potential that acts on the atoms. A section of these is shown in Figure 19, as dashed contours. The dotted contours show a fit of the form $K^2 = a_1 + a_2x^4 + a_3x^2 + a_4y^2 + a_5z^2$, where the origin is the centre of the trap. The long axis of the trap is z , and it splits along the x axis. The section is taken in the $x - y$ plane, near the potential minimum on the z axis.

23. Gas microscope

Instead of trapping atoms near a microchip, where the electromagnetic fields that manipulate them can be radiated from wires on the surface, more recent experiments by the Greiner group at Harvard held the atoms near the surface of an optical lens, where a laser beam can form a near-field image [30]. These atoms are trapped by a sheet of laser light in a plane parallel to the lens surface. Meanwhile, a pattern of evanescent waves on the surface forms an image in the plane where the atoms are trapped; the intensity of this image sets the potential across the plane. The lens was built into the wall of the vacuum chamber, and it could resolve features on the scale of the atomic separation. By changing the image over time from a single line of pixels to two lines, a line of atoms could be split in two in a similar way as with an atom chip.

The atoms in these experiments are observed by a camera that can locate single atoms in a condensed cloud [64, 29]. The principle is simple. The cloud is stored in an optical lattice, so that atoms can tunnel between adjacent cells. The depth of the lattice can be manipulated to release the atoms and capture them at a later time, a mechanism that allows the dynamics of the gas to be stopped and makes the atoms sit still while they are photographed. Meanwhile, a sophisticated lens system [29] allows a laser beam to focus on each cell of the lattice and count the number of atoms in the cell.

This system allows number correlations to be measured directly; Section 28 describes some interesting phenomena that might be observed this way.

24. Bisection as beam splitting

In the classic double slit experiment, light first shines on a single narrow slit, which diffracts a beam wide enough to illuminate both of the double slits. The light that falls on the double slits is diffracted through them, and these beams are wide enough to overlap and form interference fringes. The atomic diffraction experiments achieve the same effect by somewhat different means.

The final part of the experiment, where the atoms are released from the split trap, is very similar to diffraction through double slits. When the atoms are confined to the split trap, they have a large momentum transverse to the lines on which they are trapped, just as light passing through one of the slits develops a momentum component perpendicular to the slit. The process by which the atomic waves spread out and overlap is exactly analogous to the way light waves diffract out of the double slits. The details will be discussed in Section 31.

The first part of the atomic experiment, where the trap is split, differs from the textbook double slit experiment. It could be done the textbook way in principle: the atoms would be released from the single trap, and allowed to expand over the area covered by the double trap. Then the double trap would be switched on, some of the atoms would be captured in it, and the rest would escape. But there are many problems that prevent this from being done in practice. Cold atoms are more scarce than photons; there is enough light left after

a double slit to form visible interference fringes, but too few atoms would be recaptured for their interference fringes to stand out from the shot noise. Also, light diffracts through small angles, but the atomic cloud would expand very quickly. The atoms would be moving fast as they were recaptured, and would be unlikely to remain coherent.

Instead of separating the atomic field into two parts by diffraction, the splitting of the trap is more like splitting a beam of light with a beam splitter. In the optical analog of the split trap experiment, a light beam would split in two at a beam splitter, then the split beams would expand by diffraction until they overlapped and interfered. An optical beam splitter is in effect a tunneling barrier for photons. A photon in the input mode can be reflected into one output mode, or it can tunnel through the splitter into the other output mode. The final quantum state is an equal superposition of a photon in the two output modes. If the photons in the output modes are counted, there will be some fluctuations due to the shot noise of the tunneling process.

The splitting process resembles a beam splitter with the atomic field in place of an optical field. The main difference is that the light beam splits as it propagates through space, but the atoms stay in one place and split over time. At the time when the barrier height has increased to match the chemical potential, an atom can tunnel between the two sides of the trap; when the barrier gets higher than that, the two sides are isolated. The split trap contains many atoms, so the numbers that tunnel into each side fluctuate from shot noise like the photon numbers through a beam splitter. The statistics are somewhat different. Atoms repel each other, so we might guess that the numbers on each side will be more nearly equal than for photons, and the phase fluctuations will increase correspondingly.

The Luttinger liquid analysis treats the trap splitting exactly like an optical beam splitter, and ignores atomic repulsion. This is necessary for practical reasons; it is not possible to identify the atomic field with a Luttinger liquid until the trap has finished splitting. The justification for this assumption is that the trap splits quickly, and the two components separate entirely before the atoms have time to equalise their numbers on each side. This must be true to some extent, because the model predicts the effective temperature of the relative phase between the components. However, it is hard to make this argument quantitative, because the splitting is a continuous process, and there are no exact times that divide a single component, two components in tunnelling contact, and two separate components. Also, as will be discussed in Section 26, the atoms are not randomly distributed across the initial trap. Repulsion causes their positions to be correlated in the ground state, which reduces the repulsion energy more than it increases the zero point energy. The goal of this chapter is to analyse the splitting process more precisely, using a truncated Wigner approach.

The discussion so far has been in terms of particles. For both photons and for condensed atoms, classical waves are usually a more physically reasonable picture. In the wave picture, a beam splitter transforms waves from its bright input mode into an even superposition of its output modes, and waves from its dark input mode into an odd superposition.

The waves in an optical beam splitter are the electromagnetic field; those in the split trap are the atomic order parameter. It is pretty obvious that bright input mode for the

split trap is the ground state of the single initial trap. After the trap splits, its ground state is degenerate, because the condensate forms two independent components. In the ground state, the number of atoms in these components is fixed; the Gross-Pitaevskii equation is nonlinear, so it does not obey a superposition principle, and there is a particular way to divide the atoms that minimises the repulsion energy. However, the two components can be combined with any phase, and the condensate will still be in the ground state of the trap. The even and odd superpositions are shown in Figure 19.


The trapping potential in the Vienna experiment was designed so that the atoms remained in their ground state throughout the trapping process [68]. As the trap splits, the ground state of the initial trap evolves adiabatically to the even superposition of the two components: the two components of the split condensate are the equivalent of the output modes of a beam splitter. The first excited state of the initial trap evolves to an odd superposition of the two components. This excited state of the order parameter corresponds to the dark input to the optical beam splitter. The initial trap has cylindrical symmetry, so its first excited state is actually degenerate. However, it is obvious from parity considerations that the state that is excited along the splitting axis becomes an odd superposition of the final condensate components.

In classical field theory, the dark input mode to a beam splitter is simply dark. The quantum description in terms of particles has shot noise, and this is missing so far from our description in terms of waves. The missing part is vacuum fluctuations in the ground state of the dark input mode. In the optical case, these provide quantum noise that preserves the uncertainty principle in the output modes; virtual photons from the dark input end up in an odd superposition of the output modes while the real photons from the bright input mode end up in an even superposition. The virtual photons interfere with the real ones to produce shot noise. The same analysis applies to the atomic order parameter in the split trap.


25. Models of the Split Trap

Section 22 discussed the phenomena that are observed in interference experiments with a split atom trap, and Section 24 described the analogy with optical double slit interference that motivated these experiments. The remainder of this chapter will simulate the experiments numerically, to verify that this analogy holds quantitatively, and to test to what extent it is affected by atomic repulsion. This requires a physical model that captures enough detail of the atomic condensate to accurately describe the experimental phenomena, while remaining simple enough to be solved by the available simulation methods. This section surveys the options.


The simplest model of the split trap is to treat the order parameter of the condensed gas as a classical field, and compute its dynamics by numerically solving the Gross-Pitaevskii equation. In principle, this field could be perfectly coherent, in which case the solution would describe entirely straight interference fringes. However, as described in Section 5, a gas of bosons in one dimension does not condense into a coherent field, but has a finite





Section of the trapping potential in the tightly confined plane through the bottom of the trap. The provided potential samples are shown as solid contours, and a quartic fit as dotted contours. The contour levels are $1/x_0^2, 4/x_0^2, 9/x_0^2, \dots, n^2/x_0^2$, where $x_0 = 0.25 \mu\text{m}$ is intermediate between the widths of the initial and final potentials. A purely quadratic potential will produce evenly spaced contours.



Ground states of the quartic fit potentials. Brightness indicates amplitude, color phase. In the analogy to an optical beam splitter, this is the bright input mode. Note that the diameter of the initial cloud is around $\frac{1}{4}$, so the three-dimensional atomic density is around 150.



First excited states of the fitted potential as it splits. In the analogy to an optical beam splitter, this is the vacuum input mode.



The input and vacuum modes, corrected for atomic repulsion. This is explained in the Hartree Modes section below. The times 1 and 16 ms have been used because the degeneracy at 0 and 17 ms makes the computed modes very sensitive to the discretisation of the potential.

Figure 19. Orbitals and corrections for repulsion.

coherence length. The temperature in the experiments was high enough for this coherence length to be shorter than the length of the line of gas, so the order parameter would at least have to be averaged over a classical thermal distribution.

The classical field would adequately describe the phase of a single line of gas. However, interference results from the phase difference between two lines. When a classical field is split, the two components are perfectly coherent with each other, apart from technical noise resulting from fluctuations in the splitting process each time it is performed. A quantum field is different, because it always has shot noise due to a finite number of atoms independently falling into one component or the other. Even when thermal fluctuations are included in the classical order parameter, it can not describe the fluctuations in the relative phase when the cloud is split.

The models used in this thesis are semiclassical, as are most other theoretical treatments of the split gas experiments. This means that the law of motion of the condensate is the Gross-Pitaevskii equation, or some approximation to it. However, quantum superpositions are treated to a limited extent by generating an ensemble of order parameters, which include a random component large enough that moments of the ensemble satisfy the uncertainty principle. The details of the method, and the correspondence to quantum dynamics, were discussed in Section 15.

The order parameter could be treated as a three-dimensional field, but this is not necessary. We can save a lot of computational effort by treating the condensate as a two-dimensional field, as drawn in Figure 15, and as it appears in the experimental photographs. The atoms are always tightly confined to the plane of these diagrams. In this two-dimensional model, the equation of motion is Equation 17. The potential of the trap, in the plane where the atoms lie, has been provided by the experimentalists.

However, this two dimensional model seems more complicated than necessary. Section 24 explained how the splitting of the atoms can be explained by adiabatic evolution of the transverse orbitals of the trap. It ought to be possible to expand the order parameter over these orbitals, and treat the components of the condensate as a set of one-dimensional fields. There are some technical complications to doing this, but the attempts are discussed below in Sections 32 and 33.

Finally, there is a simpler model for condensate dynamics than the Gross-Pitaevskii equation. For small excitations such as shot noise, the dynamics of a condensate can be treated as sound waves in the order parameter. Some existing work [8] has applied this approach to the one-dimensional fields after the condensate is split. This has the limitation that it can not treat the splitting process itself, but it gives a simple picture of what happens afterwards, when the two components of the condensate evolve independently. The sound wave model of a condensate is also technically useful in setting up the ensemble of order parameters that represent the state of the condensate.

In order to use any of these models numerically, the continuous atomic field must be expanded over a discrete set of orbitals. The simulation method used is most accurate when there are as few orbitals as possible, and each of them is occupied by many atoms. This is

the classical wave limit; presumably all semiclassical methods are most accurate in these conditions.

The orbitals are defined by grid points (x_i, z_j) , where the spacing of the x_i is l and that of the z_j is h . When the atomic field is projected onto the low-momentum space defined by this grid, it can be written as an expansion

$$\hat{\psi}(x, z) = \frac{1}{\sqrt{lh}} \sum_{ij} \hat{a}_{ij} \operatorname{sinc}\left(\frac{x - x_i}{l}\right) \operatorname{sinc}\left(\frac{z - z_j}{h}\right). \quad (38)$$

The \hat{a}_{ij} are annihilation operators with $[\hat{a}_{ij}, \hat{a}_{kl}^\dagger] = \delta_{ik}\delta_{jl}$, so that the $\hat{\psi}$ have a momentum-limited commutation relation

$$[\hat{\psi}(x, z), \hat{\psi}^\dagger(x', z')] = \frac{1}{lh} \operatorname{sinc}\left(\frac{x - x'}{l}\right) \operatorname{sinc}\left(\frac{z - z'}{h}\right).$$

Note that the values $\operatorname{sinc}(x - x_i)/l$ can be interpreted as samples of the function $f(x') = \operatorname{sinc}(x - x')/l$ at the points x_i . The function f is bandwidth limited, therefore

$$f(x') = \sum_i f(x_i) \frac{\operatorname{sinc}(x' - x_i)}{l} = \sum_i \frac{\operatorname{sinc}(x' - x_i)}{l} \frac{\operatorname{sinc}(x - x_i)}{l} = \operatorname{sinc}(x - x')/l.$$

It will be convenient to define interpolated operators

$$\hat{a}_i(z) = \sum_j \hat{a}_{ij} \operatorname{sinc}\left(\frac{z - z_j}{h}\right),$$

so that

$$\hat{\psi}(x, z) = \frac{1}{\sqrt{lh}} \sum_i \hat{a}_i(z) \operatorname{sinc}\left(\frac{x - x_i}{l}\right). \quad (39)$$

These have the commutation relation for discrete orbitals, $[\hat{a}_i(z), \hat{a}_j^\dagger(z)] = \delta_{ij}$.

26. Bogoliubov ground states

So far, this thesis has focused on coherent states of quantum fields. These are often a good description for the electromagnetic field; photons hardly interact, so they distribute themselves among field modes almost at random. This results in a Poisson distribution of photon numbers, which is precisely what a multimode coherent state describes. An atomic field, on the other hand, becomes coherent only when the forces between the atoms are small.

The exact ground state is known for a one dimensional field of bosons that repel each other when they collide; it is called the Lieb-Liniger solution [5]. The multiparticle wave function for the ground state is composed piecewise from wave functions $\psi(x_1, \dots, x_n) = \exp i(k_1 x_1 + \dots + k_n x_n)$ for free particles, with appropriate boundary conditions on the hyperplanes $x_i = x_j$ where two particles collide. In order for this state to be stationary, it turns out that the set of momenta $\{k_1, \dots, k_n\}$ must be the same in every piece; the particles

can be distinguished only by their order along the line, and there are different probability amplitudes for the various orders in which these momenta can be placed.

The Lieb-Liniger state has a reasonably simple form in first quantisation, except that the combinatorially large number of particle orders makes computation impractical. It is a very complicated form in second quantisation, where the particles are distributed among a fixed set of orbitals. It is not practical to derive a Wigner distribution for this state or to generate samples, although it does allow some exact correlation functions to be derived that are useful for validating numerical approximations [69]. Fortunately, an approximate ground state was derived by Holstein and Primakoff (on one side of the iron curtain) [70], and independently by Bogoliubov (on the other side) [67]. This becomes exact in the classical limit where the Lieb-Liniger parameter γ is small, and it has a simple form in second quantisation.

When repulsion is significant, the potential energy is minimised when the atoms are regularly spaced. On the other hand, regular spacing increases the zero point energy required to constrain their positions. The healing length is the width of the smallest hole that the rest of the condensate can make for each atom, where the zero-point pressure required to form the hole balances the repulsion pressure of the bulk condensate. In the Tonks gas regime, this length is shorter than the atomic separation, allowing each atom to stay in its own block without overlapping the rest of the condensate.

When the healing length is on the same scale as the atomic separation, the atoms must overlap, but they will be spaced more regularly than in a coherent state, causing the mode occupation numbers to be better defined than in a coherent state. The reduction in number fluctuations caused by atomic repulsion will increase the phase fluctuations, and all this will affect the coherence that we are studying. In the split trap, these effects are large enough to matter, and the coherent state does not adequately approximate the ground state of the gas.

It might be apparent to readers from a quantum optics background that this variation in the size of the number fluctuations sounds somewhat like squeezing. This section derives the Bogoliubov transformation, which maps a vacuum state to an approximate ground state; it turns out to be exactly a squeezing transformation. The Wigner functions for squeezed states are positive, and this state can be sampled for Monte-Carlo purposes.

The Bogoliubov approximation is valid when the state of the atomic field is close to a coherent state with the equilibrium order parameter $\psi(x)$ as its amplitude. The repulsion term in the Hamiltonian is fourth order in the field operators $\hat{\Psi}(x)$ and $\hat{\Psi}^\dagger(x)$. The assumption means that the operator $\delta\hat{\Psi} = \hat{\Psi} - \psi$ is small, and the terms that are cubic or quartic in it can be neglected, leaving an approximate Hamiltonian that is quadratic in the field $\delta\hat{\Psi}$. The ground state of this Hamiltonian should approximate the ground state of the full Hamiltonian better than the coherent state does.

The equilibrium order parameter defines an orbital, and by assumption the state of this condensate orbital is coherent. The ground state of the other orbitals remains to be determined. Those orbitals can not have a nonzero coherent amplitude; if they did, it could

be absorbed into the order parameter. So there are no linear terms in the approximate Hamiltonian, and its minimum will lie at zero coherent amplitude. This is shown more rigorously below, where the Bogoliubov transformation is derived. The remaining question is the shape of the quadratic energy well. It is possible for this to be distorted in ways that make no difference to the classical equilibrium point, but cause the ground state of the non-condensate orbitals to be a squeezed state instead of a vacuum.

The goal is to find harmonic modes whose amplitudes are independent, with amplitudes a_i that split the Hamiltonian of Equation 14 into a sum of the form $\sum_i \chi_i a_i^\dagger a_i$ with no cross terms. Because the Hamiltonian is nonlinear in the field amplitude, these modes might not correspond to orbitals in a simple way. In particular, the coordinate $a_i + a_i^\dagger$ is the component of some orbital ϕ_i in the field ψ , and the canonical momentum $i(a_i^\dagger - a_i)$ is the component of some orbital ϕ'_i in $\frac{d\psi}{dt}$, but ϕ and ϕ' are not necessarily the same orbital.

There are two ways to determine the harmonic modes, and therefore the form of the quadratic approximation to the Hamiltonian. The first is to diagonalise the Hamiltonian by brute force, applying some squeezing transformations that give a set of independent harmonic oscillators. This is the Bogoliubov transformation, described later in this section. The tricky way is to realise that, when the classical field is perturbed by adding a component along one of the harmonic modes, the trajectory of the perturbed field will be harmonic. A search for harmonic solutions to the Gross-Pitaevskii equation, in the limit that the classical field is close to the equilibrium point, will reveal the modes that diagonalise the Hamiltonian. We will follow this classical approach first, then introduce the quantum version.

The classical order parameter field can be interpreted as a description of a classical fluid; the squared modulus of the order parameter is the fluid density, and the gradient of its phase is the mass flux. The harmonic solutions to the Gross-Pitaevskii equation correspond to sound waves in this fluid, where the quantum gas is rarefied and compressed. The waves do not obey a superposition principle in general, because the GPE is nonlinear, but they do so in the limit where their amplitudes are small. In this limit, the sound waves can be quantised by treating the modes as independent oscillators; the imposition of this superposition principle is how the Bogoliubov Hamiltonian differs from the original atomic Hamiltonian. Hopefully, the vacuum fluctuations in the ground state of the field will be small enough for this approximation to be accurate. This assumption will be examined after the ground state is derived.

It remains to find the mode functions, and the harmonic frequencies for sound waves in the condensate fluid. The derivation follows [1], using the scaled units of Section 3. When the Gross-Pitaevskii equation has the form of Equation 15, the equilibrium order parameter satisfies the eigenproblem

$$(-\nabla^2 + K^2 + \xi^{-2})\psi_0 = \Xi^{-2}\psi_0, \quad (40)$$

and the order parameter evolves in time as

$$\psi(s, x) = \psi_0(x) e^{-\frac{i}{2\Xi^2}s}.$$

Here Ξ is the length derived from the chemical potential using dispersion units, which is often close to the healing length in the densest part of the condensate.

We start by subtracting $\Xi^{-2}\psi$ from the right hand side of Equation 15. This is just a displacement of the energy scale, and has no physical effect; however, it causes the chemical potential to be zero in equilibrium, so that a sound wave with physical frequency χ is represented by a field with a static component and a component whose phase rotates at frequency χ , instead of components at Ξ^{-2} and $\Xi^{-2} + \chi$. This gives a classical equation of motion

$$2i\partial_s\psi = (-\nabla^2 + K^2 + g|\psi|^2 - \Xi^{-2})\psi. \quad (41)$$

We seek sound wave solutions of the form $\psi = \psi_0 + \delta\psi$, where ψ_0 is constant and $\delta\psi$ is harmonic. When this form is substituted in Equation 41, the left hand side splits into a constant term and a term proportional to $\delta\psi$; on the right hand side, the expansion of the term in $|\psi|^2\psi$ includes a term in $\psi_0^2\delta\psi^*$ whose phase rotates in the opposite sense to the phase of $\delta\psi$. Therefore $\delta\psi$ must rotate in both senses, and the solution has the form

$$\psi = \psi_0 + ue^{-i\chi s} - v^*e^{i\chi s}. \quad (42)$$

The solutions come in identical pairs due to the symmetry where χ is negated and u is swapped with $-v^*$, but we can select one of each pair by requiring the frequency χ to be positive.

When the form of Equation 42 is substituted, Equation 41 becomes

$$2\chi(ue^{-i\chi s} + v^*e^{i\chi s}) = (-\nabla^2 + K^2 - \Xi^{-2})\delta\psi - g|\psi_0|^2\psi_0 + g|\psi|^2\psi.$$

The last term is a polynomial that is quadratic in ψ_0 and $\delta\psi$, and linear in their complex conjugates. When it is expanded, the static terms $g|\psi_0|^2\psi_0$ cancel, and every remaining term is at least linear in $\delta\psi$ or its conjugate. Collecting the linear terms that rotate as $e^{-i\chi s}$ and $e^{i\chi s}$ gives an eigenproblem called the Bogoliubov equations [71],

$$\begin{aligned} (-\nabla^2 + K^2 + 2g|\psi_0|^2 - \Xi^{-2})u - g\psi_0^2v &= 2\chi u, \\ (-\nabla^2 + K^2 + 2g|\psi_0|^2 - \Xi^{-2})v - g\psi_0^2u &= -2\chi v. \end{aligned} \quad (43)$$

Sound waves that satisfy this eigenproblem satisfy the Gross-Pitaevskii equation to first order in $u/|\psi_0|$ and $v/|\psi_0|$. This is adequate when the field is nearly coherent, so that u , v and $\delta\psi$ are small.

The obvious way to compute solutions to the Bogoliubov equations is to transform them into the eigenvalue problem

$$\begin{pmatrix} -\nabla^2 + K^2 + 2g|\psi_0|^2 - \Xi^{-2} & -g\psi_0^2 \\ g\psi_0^2 & \nabla^2 - K^2 - 2g|\psi_0|^2 + \Xi^{-2} \end{pmatrix} \begin{pmatrix} u_i \\ v_i \end{pmatrix} = 2\chi \begin{pmatrix} u_i \\ v_i \end{pmatrix}, \quad (44)$$

which can be solved by many well-known algorithms. However, this turns out to be numerically unstable. When the frequencies of the modes are degenerate, the degenerate mode functions that are computed can be nearly identical, which results in nearly equal quantities

being subtracted when the quantum ground state is computed. The problems are resolved [72] by solving for the functions $\Psi^+ = u + v$ and $\psi^- = u - v$, which satisfy the independent equations

$$\begin{aligned}(h_0^2 + 2g|\psi_0|^2 h_0)\psi^+ &= 2\chi^2\psi^+ \\ (h_0^2 + 2gh_0|\psi_0|^2)\psi^- &= 2\chi^2\psi^-, \end{aligned} \tag{45}$$

where $h_0 = -\nabla^2 + K^2 + g|\psi_0|^2 - \Xi^{-2}$. The sum and difference modes are related by $\psi^- = (h_0/2\chi)\psi^+$. It is not obvious why this problem is more stable—the operators being diagonalised are not normal—but in practice it is. When the modes are degenerate, it is best to solve Equations 45 for ψ^+ and use this relation to find the corresponding ψ^- . When both ψ^+ and ψ^- are computed from Equations 45, it is not obvious how to match up each ψ^- eigenvector with the particular linear combination of the degenerate ψ^- eigenvectors that corresponds to it.

For a given potential, there will be a discrete set of modes u_i and v_i that satisfy Equations 43. These do not satisfy any orthogonality relation, so it is not possible to expand an arbitrary $\delta\psi$ in the usual way, using inner products. However, given a perturbation $\delta\psi$, the expansion $\delta\psi = \sum_i u_i\alpha_i - v_i^*\alpha_i^*$ is a linear combination of certain modes related to u_i and v_i , in which the coefficients are the real and imaginary parts of α_i . This can be solved, possibly in a least squares sense, to give an expansion of $\delta\psi$ over the sound wave modes.

When ψ_0 is real, as an equilibrium order parameter usually is, Equations 43 have real solutions for u and v [73]. However, the form of Equation 42 is a complex field. So in what sense does the harmonic dynamics generated by u and v describe a physical sound wave?

At times $\chi s = n\pi$ with n an integer, the terms $ue^{-i\chi s}$ and $v^*e^{i\chi s}$ are real, and they are small perturbations of the large real amplitude ψ_0 . At these times, the order parameter is stationary, and the fluctuation $\delta\psi = \pm(u - v^*)$ perturbs the density of the condensate. The compressions and rarefactions swap places every half cycle. This is what a standing sound wave would do.

Between these times, when $\chi s = (n + \frac{1}{2})\pi$, the quantity $\delta\psi = \pm i(u + v^*)$ is imaginary. To first order, a purely imaginary perturbation to the real quantity ψ_0 has no effect on the modulus of the order parameter: these are the times when the standing wave passes through the equilibrium density $|\psi_0|^2$. The imaginary quantity $\delta\psi$ twists the phase of the order parameter, giving the condensate a velocity proportional to $\nabla(u + v^*)/|\psi_0|$. The condensate is a nonlinear fluid that supports unusual sound waves whose density and velocity peaks have different forms. However, this is only the case at long wavelengths. In the limit where the wavelength is much shorter than the healing length, the nonlinear repulsion term becomes small compared to the kinetic energy, and the modes revert to ordinary linear sound waves, with u much larger than v^* .

So far, we have analysed sound waves in the classical order parameter. The quantum mechanical analysis treats the atomic field $\hat{\Psi}$, which differs from the order parameter ψ_0 by a small perturbation $\delta\hat{\Psi}$, with

$$\hat{\Psi}(x) = \psi_0(x) + \delta\hat{\Psi}(x). \tag{46}$$

The modes of Equation 42 become an expansion known as the Bogoliubov transformation [67]

$$\delta\hat{\Psi}(x) = \sum_i u_i(x)a_i - v_i^*(x)a_i^\dagger \quad (47)$$

that relates the field operator to a set of phonon mode operators a_i . Physically, this means that a sound wave is excited when the creation operator adds atoms to the compressed parts of the fluid and the annihilation operator removes them from the rarefied parts. Exciting a sound wave should not change the number of atoms in the gas, so in a strict analysis the atoms created and annihilated would move to and from the condensed mode a_0 , and $\delta\hat{\Psi}$ would include terms in $u_i a_i a_0^\dagger$ and $v_i^* a_i^\dagger a_0$. However, the state of the condensed mode is assumed to be coherent with a large amplitude, so that $a_0|\Psi\rangle \approx a_0^\dagger|\Psi\rangle \approx \sqrt{N}|\Psi\rangle$, and the factor of \sqrt{N} related to the number of atoms is absorbed into u_i and v_i .

Substituting Equation 46 into the atomic field Hamiltonian of Equation 14, and collecting terms by order of the perturbation $\delta\hat{\Psi}$, gives [1]

$$\begin{aligned} H = & \int -\psi_0^* \nabla^2 \psi_0 + K^2 |\psi_0|^2 + \frac{g}{2} |\psi_0|^4 \\ & + \int (-\nabla^2 + K^2 \psi_0 + g |\psi_0|^2) \psi_0 \delta\hat{\Psi} + \text{h.c.} \\ & + \int -\delta\hat{\Psi}^\dagger \nabla^2 \delta\hat{\Psi} + K^2 \delta\hat{\Psi}^\dagger \delta\hat{\Psi} + \frac{g}{2} \left((\delta\hat{\Psi}^\dagger)^2 \psi_0^2 + \delta\hat{\Psi}^2 (\psi_0^*)^2 + 4\delta\hat{\Psi}^\dagger \delta\hat{\Psi} |\psi_0^*|^2 \right) \\ & + O(\delta\hat{\Psi}^3). \end{aligned}$$

The first line is the chemical potential Ξ^{-2} of the equilibrium order parameter. It can be set to zero by referring the external potential K^2 to Ξ^{-2} , and we will assume this has been done. The second line is proportional to the right-hand side of the Gross-Pitaevskii equation, so, as was claimed above when the Bogoliubov approximation was introduced, the Hamiltonian is minimised to first order in $\delta\hat{\Psi}$ when the order parameter assumes its equilibrium value.

The third line is second order in $\delta\hat{\Psi}$. It can be shown [73] that the eigenvalue property of Equation 43 reduces the spatial integrals to sums over such terms as $a_i^\dagger a_j^\dagger \int u_i^* v_j$ and $a_i^\dagger a_j \int u_i^* u_j$. At this point, we will assume that the mode functions are real; they have completeness and orthogonality properties that diagonalise these sums. The Hamiltonian reduces to

$$H = \int -\psi_0^* \nabla^2 \psi_0 + K^2 |\psi_0|^2 + \frac{g}{2} |\psi_0|^4 - 2 \sum_{i \neq 0} \chi_i \left(a_i^\dagger a_i + \int |v_i|^2 \right) \quad (48),$$

where χ_i is the eigenvalue from Equation 43 for the modes u_i and v_i . The orthogonality properties are derived from the requirement that the $\delta\hat{\Psi}$ of Equation 47 satisfies an appropriate field operator commutation relation whenever a_i satisfies $[a_i, a_i^\dagger] = 1$. They are not strictly orthogonality requirements, because they can not be derived from an inner product on the u_i and v_i . If we define $w_i = (u_i, v_i)$ for convenience, the two requirements [73] are

$$\mathcal{F}[w_j, w_k] = \int u_j u_k - v_j v_k = \delta_{jk} \quad \text{and} \quad \mathcal{G}[w_j, w_k] = \int u_j v_k - v_j u_k = 0, \quad (49)$$

These are not quite an orthogonality relation: the functionals \mathcal{F} and \mathcal{G} are bilinear, but they are not positive definite. Also, \mathcal{F} is symmetric, and \mathcal{G} is antisymmetric. This implies a normalisation condition

$$\mathcal{F}[w_i, w_i] = \int |u_i|^2 - |v_i|^2 = 1. \quad (50)$$

We restricted the solutions of Equation 43 to those with positive χ . These can be shown to form a complete set [73]. Given this, the ground state of the Hamiltonian of Equation 48 is obvious: it occurs when all the phonon modes a_i are in their vacuum states. In the truncated Wigner model, the ensemble of classical fields is found from Equation 47, treating $\hat{\Psi}$ as a classical field instead of a quantum one. Each mode operator a_i becomes a Gaussian distributed classical amplitude α_i as usual. When the Bogoliubov transformation of Equation 47 is applied, the real part of α_i is shrunk by a factor $u - v$, and the imaginary part stretched by a factor $u + v$. This resulting distribution is still Gaussian, but it forms an ellipse instead of a circle in the complex plane.

It remains to find the modes u and v [1]. In an untrapped condensate in one dimension, the constant potential may be set to $K^2 = 0$, in which case the chemical potential of Equation 40 is $\Xi^{-2} = \xi^{-2} = g|\psi_0|^2$. The eigenmodes of Equations 43 are proportional to the eigenfunctions of the Laplacian; since we are assuming the modes to be real, these are $\sin(kx)$ and $\cos(kx)$. These can be chosen with k positive. Physically, these describe standing sound waves, which translate more neatly to the modes of a trapped gas. When such sinusoidal modes are substituted into Equations 43, the result is a 2×2 matrix eigenproblem,

$$\begin{pmatrix} k^2 + \xi^{-2} & -\xi^{-2} \\ \xi^{-2} & -k^2 - \xi^{-2} \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = 2\chi \begin{pmatrix} u \\ v \end{pmatrix}. \quad (51)$$

In this equation, the symbols u and v can be read as sinusoidal functions; alternatively, they can be read as the amplitudes. The solution is $2\chi = k\sqrt{k^2 + 2\xi^{-2}}$. In the limit where the wavelength is long compared to the healing length, so that $\sqrt{k^2 + 2\xi^{-2}} \approx \sqrt{2}\xi^{-1}$, this asymptotes to a linear dispersion relation, with the frequency χ proportional to k . In this limit, the modes are very much like sound waves, and propagate at the Bogoliubov sound speed $1/(\sqrt{2}\xi)$.

For a more precise analysis, let us impose periodic boundary conditions in a one-dimensional gas of length L , and treat u and v as the amplitude coefficients of normalised mode functions, with pairs of modes

$$u(x) = u\sqrt{2/L} \sin kx, \quad v(x) = v\sqrt{2/L} \sin kx,$$

and

$$u(x) = u\sqrt{2/L} \cos kx, \quad v(x) = v\sqrt{2/L} \cos kx \quad (52)$$

whose energies are degenerate. In this case, Equation 50 becomes $u^2 - v^2 = 1$. Cross-multiplying Equations 51 gives

$$(k^2 + \xi^{-2} - 2\chi)u^2 = (k^2 + \xi^{-2} + 2\chi)v^2,$$

whence

$$u^2 = \frac{\kappa + \kappa^{-1}}{2\sqrt{\kappa^2 + 2}} + \frac{1}{2} \quad \text{and} \quad v^2 = \frac{\kappa + \kappa^{-1}}{2\sqrt{\kappa^2 + 2}} - \frac{1}{2}, \quad (53)$$

where $\kappa = k\xi$ is the wave number referred to the healing length. The results for u , v and χ are shown in Figure 20, along with ellipses to indicate the squeezing of a Bogoliubov mode at each wave number.

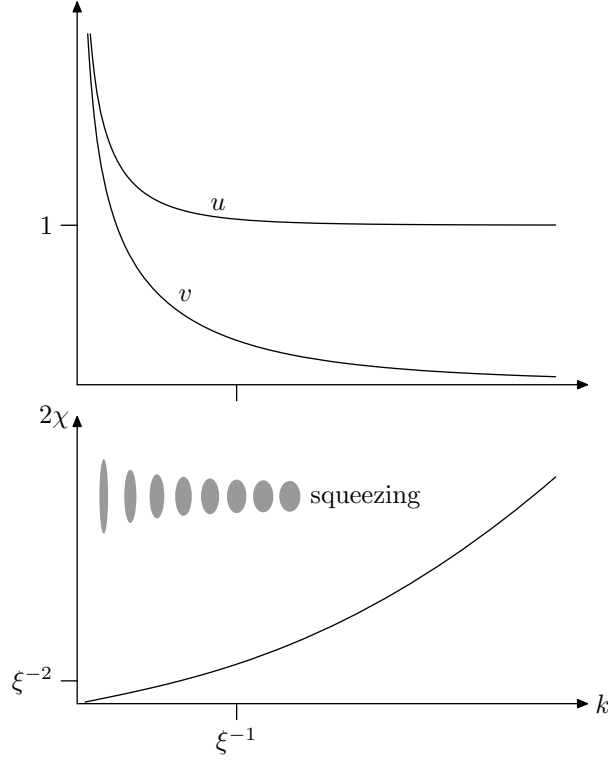


Figure 20. Amplitudes and energies of Bogoliubov modes.

In the high energy limit $\kappa \rightarrow \infty$, the asymptotes are $u^2 \rightarrow 1$ and $v^2 \rightarrow 1/4\kappa^4$. In this limit, the sound wave modes become identical to standing wave orbitals: the Bogoliubov transformation has no effect, and Equation 47 reduces to an expansion of $\delta\hat{\Psi}$ over these orbitals. This means that exciting a sound wave with a given wave number is identical to exciting an atom from the condensate to an orbital with the same wave number; the quantisation of the sound waves is due to the discrete nature of the atomic gas.

In the continuous limit $\kappa \rightarrow 0$, both u^2 and v^2 asymptote to $1/(2\sqrt{2}\kappa)$. However, the longest mode in a gas of length L has wave number $k_0 = 2\pi/L$, where, for a gas of N particles, $L = \sqrt{2\gamma}N\xi$. Assuming that κ is small, which is the case whenever this mode can become large, its norm is

$$u^2 \approx \frac{\sqrt{\gamma}}{4\pi}N,$$

where γ is the Lieb-Liniger parameter defined in Section 8. The mean field approximation requires this mode to be small compared to the order parameter of N particles, which is

true only in the classical case of small γ .

In a trap, the mode shapes are determined by the condensate density $|\psi_0|^2$. The wave numbers are discrete, and the condensate size will set a separation between them. The sound speed decreases to zero in the limit of low atomic density, and the boundary conditions will depend on the shape of the trap at the edges of the condensate. In most cases, the modes will expand and contract somewhat at the trap edges. The lowest frequency modes of a trapped condensate are not like sound waves at all, but are Kohn and breathing modes of the whole condensate.

In symmetric traps, and particularly in the case of an untrapped gas, several sound wave modes can have the same frequency. Such degenerate modes satisfy Equation 44, but do not necessarily have the orthogonality properties that generate the diagonal form of Equation 48. Numerical diagonalisation often computes modes of this kind, and the orthogonality properties of Equation 49 must be imposed explicitly by taking appropriate linear combinations of the computed eigenvectors.

Suppose that $w_1 = (u_1, v_1)$ and $w_2 = (u_2, v_2)$ are a pair of degenerate sound wave modes, part of a computed solution to Equations 45. The computed modes will not necessarily satisfy Equations 49, so these constraints have to be imposed explicitly. The obvious method is to look for linear combinations $w_a = w_1 + \lambda w_2$ and $w_b = w_2 + \mu w_1$ that satisfy Equations 49 apart from the normalisation condition, then normalise these.

One requirement is

$$\mathcal{G}[w_a, w_b] = \mathcal{G}[w_1 + \lambda w_2, w_2 + \mu w_1] = (1 - \mu\lambda)\mathcal{G}[w_1, w_2] = 0.$$

This can be satisfied in two ways: either the modes w_1 and w_2 satisfy the condition already, with $\mathcal{G}[w_1, w_2] = 0$, or else $\mu\lambda = 1$. The other requirement is

$$\mathcal{F}[w_a, w_b] = (1 + \mu\lambda)\mathcal{F}[w_1, w_2] + \lambda\|w_2\|^2 + \mu\|w_1\|^2 = 0. \quad (54)$$

Here $\mathcal{F}[w, w]$ is written $\|w\|^2$, although this is not actually a norm. In the case $\mathcal{G}[w_1, w_2] = 0$, this can be solved by setting $\mu = 0$ and $\lambda = -\mathcal{F}[w_1, w_2]/\|w_2\|^2$, or alternatively by $\lambda = 0$ and $\mu = -\mathcal{F}[w_1, w_2]/\|w_1\|^2$. The modes are degenerate, so λ and μ are not necessarily unique. The other case, where $\mu\lambda = 1$, guarantees that $\lambda \neq 0$, so Equation 54 can be reduced to

$$\lambda^2\|w_2\|^2 + 2\lambda\mathcal{F}[w_1, w_2] + \|w_1\|^2 = 0,$$

with solution

$$\lambda = \frac{-X \pm \sqrt{X^2 - \|w_1\|^2\|w_2\|^2}}{\|w_2\|^2}$$

and

$$\mu = \frac{-X \pm \sqrt{X^2 - \|w_1\|^2\|w_2\|^2}}{\|w_1\|^2},$$

where $X = \mathcal{F}[w_1, w_2]$.

27. Number fluctuation observables as Wigner distributions

The Harvard group can directly measure the difference in atomic density between the two split clouds [64]. It is convenient to measure these fluctuations using a correlation observable $:(N_l - N_r)^2:$, where N_l and N_r are the numbers of atoms in the two components after the trap is split, and the colons denote normal ordering. The expectation value of this observable is zero when the two components are coherent fields with the same amplitude; it will be negative when their numbers are correlated and positive when their phases are correlated.

In the two-dimensional description of the split trap, the atom numbers are related to the field operators as

$$N_l(z) = \int_0^\infty dx \int_{-\infty}^\infty dz \hat{\psi}(x, z)^\dagger \hat{\psi}(x, z) \quad \text{and} \quad N_r(z) = \int_{-\infty}^0 dx \int_{-\infty}^\infty dz \hat{\psi}(x, z)^\dagger \hat{\psi}(x, z).$$

Substituting the mode expansion from Section 25, it follows that

$$\begin{aligned} N_l(z) &= \frac{1}{l\hbar} \sum_{ij} \int_0^\infty dx \int_{-\infty}^\infty dz a_i^\dagger(z) a_j(z) \operatorname{sinc}\left(\frac{x-x_i}{l}\right) \operatorname{sinc}\left(\frac{x-x_j}{l}\right) \\ &\approx \sum_{i,k:x_i>0} a_{ik}^\dagger a_{ik}. \end{aligned}$$

The approximation here is that $l^{-1} \int_0^\infty \operatorname{sinc}\left(\frac{x-x_i}{l}\right) \operatorname{sinc}\left(\frac{x-x_j}{l}\right) = \delta_{ij}$. This is a good approximation for points distant from the trap axis, where the integral can be extended to $-\infty$. When x_i and x_j are small, there are discrepancies in the coefficients of $a_i^\dagger(z) \hat{a}_j(z)$ that are around 10%. This does not matter after the trap has been bisected, but it will affect the computed observables beforehand, when there is a single condensate trapped along the z axis and the atoms are being counted on either side of the axis. A wavenumber-limited field does not naturally divide into two parts sharply divided by an axis. It would be straightforward to include the cross terms if necessary, but the number differences at this point of the experiment are not very interesting, and these are omitted for simplicity.

The task now is to put the correlation observable into symmetric form in a_i and a_i^\dagger , so that it can be evaluated as an average over samples from a Wigner distribution. The cross term is straightforward:

$$\begin{aligned} :N_l N_r: &= :N_l: :N_r: \\ &= \sum_{i:x_i>0,j} a_{ij}^\dagger a_{ij} \sum_{k:x_k<0,l} a_{kl}^\dagger a_{kl} \\ &= \sum_{i:x_i>0,j} \left(s(a_{ij}^\dagger a_{ij}) - \frac{1}{2} \right) \sum_{k:x_k<0,l} \left(s(a_{kl}^\dagger a_{kl}) - \frac{1}{2} \right). \end{aligned}$$

Here s is the symmetric ordering defined in Section 15. Suppose the random complex field ψ is distributed according to the Wigner distribution $\mathcal{W}^{-1} \hat{\Psi}$. We will write this as $\psi \leftrightarrow \hat{\Psi}$. It follows from Equation 39 that symmetric moments of a_{ij} and a_{ij}^\dagger correspond to

moments as $a_{ij} \leftrightarrow \sqrt{lh} \psi(x_i, z_j)$, whence

$$:N_l N_r: \leftrightarrow l^2 h^2 \sum_{i:x_i > 0, j} \left(|\psi(x_i, z_j)|^2 - \frac{1}{2lh} \right) \sum_{k:x_k < 0, l} \left(|\psi(x_k, z_l)|^2 - \frac{1}{2lh} \right).$$

Note that $lh \sum_{ij}$ is the way that the XPDE software evaluates an integral; this software is used below to compute these observables.

The self terms are slightly more complicated.

$$:N_l^2(z): = \sum_{i,j,k,l:x_i,k > 0} a_{ij}^\dagger a_{kl}^\dagger a_{ij} a_{kl}.$$

This sum includes diagonal terms which require the identity

$$:a^{2\dagger} a^2: = s(a^{2\dagger} a^2) - 2s(a^\dagger a) + \frac{1}{2}. \quad (55)$$

The diagonal and off-diagonal cases evaluate to

$$\begin{aligned} :N_l^2(z): &= \sum_{i,j,k,l:x_i,k > 0} \begin{cases} s(a_{ij}^\dagger a_{kl}^\dagger a_{ij} a_{kl}) - s(a_{ij}^\dagger a_{ij}) - s(a_{kl}^\dagger a_{kl}) + \frac{1}{2} & i = j \text{ and } k = l, \\ (s(a_{ij}^\dagger a_{ij}) - \frac{1}{2})(s(a_{kl}^\dagger a_{kl}) - \frac{1}{2}) & \text{otherwise.} \end{cases} \\ &\leftrightarrow l^2 h^2 \sum_{i,j:x_i > 0} \left(|\psi(x_i, z_j)|^2 - \frac{1}{2lh} \right) \sum_{k,l:x_k > 0} \left(|\psi(x_k, z_l)|^2 - \frac{1}{2lh} \right) \\ &\quad - lh \sum_{i,j:x_i > 0} \left(|\psi(x_i, z_j)|^2 - \frac{1}{4lh} \right). \end{aligned}$$

The term $:N_r^2(z):$ can be evaluated similarly; the result differs only in that the sums are taken over $x_i < 0$ instead of $x_i > 0$.

Putting the terms back together, and making use of the function sgn that takes the value 1 for positive numbers and -1 for negative ones, gives

$$:(N_l - N_r)^2: \leftrightarrow \left(hl \sum_{ij} \text{sgn}(x_i) \left(|\psi(x_i, z_j)|^2 - \frac{1}{2hl} \right) \right)^2 + hl \sum_{ij} \left(-|\psi(x_i, z_j)|^2 + \frac{1}{4hl} \right). \quad (56)$$

As a simple case for tests and validation, Section 28 will consider correlations in a one-dimensional untrapped gas. These correlations will be measured by the second-order correlation function $g^{(2)}$. The value of this function is known both for the exact interacting ground and thermal states, and also for the approximate Bogoliubov states [69]; this makes it convenient for numerical validation.

The function, in the normalised form that we will use, is defined by

$$g^{(2)}(x, y) = \frac{\langle \hat{\Psi}^\dagger(x) \hat{\Psi}^\dagger(y) \hat{\Psi}(y) \hat{\Psi}(x) \rangle}{\langle \hat{\Psi}^\dagger(x) \hat{\Psi}(x) \rangle \langle \hat{\Psi}^\dagger(y) \hat{\Psi}(y) \rangle}. \quad (57)$$

For simplicity of programming, we will mostly consider the single point limit $x = y$. The denominator in this case is simply a density, and when $x = x_i$, one of the points at which the field is numerically sampled in Equation 38, it can be Wigner sampled as

$$\hat{\Psi}^\dagger(x_i) \hat{\Psi}(x_i) = \frac{1}{h} a_i^\dagger a_i \leftrightarrow |\psi(x_i)|^2 - \frac{1}{2h}. \quad (58)$$

The numerator follows from the identity in Equation 55,

$$(\hat{\Psi}^\dagger)^2(x_i)\hat{\Psi}^2(x_i) \leftrightarrow |\psi(x_i)|^4 - \frac{2}{h}|\psi(x_i)|^4 + \frac{1}{2h^2}. \quad (59)$$

The boson density can be taken in momentum space as well as position space, with

$$n(k_i) = \hat{\Psi}^\dagger(k_i)\hat{\Psi}(k_i) \leftrightarrow |\tilde{\psi}(k_i)|^2 - \frac{1}{2h_k}, \quad (60)$$

where $\tilde{\psi}(k)$ and $\psi(x)$ are related by the Fourier transform of Equation 26, and h_k is the spacing between the points k_i that form the reciprocal grid to x_i . The kinetic energy density $k^2n(k)$ can be evaluated similarly as

$$k_i^2\hat{\Psi}^\dagger(k_i)\hat{\Psi}(k_i) \leftrightarrow k_i^2 \left(|\tilde{\psi}(k_i)|^2 - \frac{1}{2h_k} \right), \quad (61)$$

and the total kinetic energy is $T = \int k^2n(k) dk$.

28. Oscillating correlations from a coherent state

Section 26 derived the Bogoliubov transformation, linking boson orbitals to sound wave and quasiparticle modes. This resolves the Hamiltonian of a field of repulsive bosons into a set of independent sound wave modes, on the assumption that the field is sufficiently coherent to neglect operators of third and higher orders in the difference $\delta\hat{\psi}$ between the boson field and its order parameter. When this assumption holds, Equation 48, along with the choice of positive frequencies χ_i , shows that the vacuum state of the sound wave modes has less energy than the coherent state.

This is a squeezing transformation. The vacuum state of the sound wave modes, which is the ground state of the approximate Hamiltonian formed by expanding the field Hamiltonian to second order in $\delta\hat{\psi}$, is a squeezed state of the boson orbitals. Conversely, a coherent state of the Bose field is a squeezed vacuum state of the sound wave modes, which is an excited state of the approximate Hamiltonian.

What would happen if condensed bosons were prepared in a coherent state, instead of their ground state? In principle, there are ways to do this. Bosonic atoms could be cooled in the presence of a magnetic field turned to a Feshbach resonance [1], a condition where their interactions are small. After these non-interacting atoms relaxed into their coherent ground state, the field could be rapidly changed, so that the atoms remained in a coherent state while they developed significant repulsion.

Atoms prepared in this way could reduce their repulsion energy by moving to a more regular spatial distribution, a state that better approximates their Bogoliubov ground state. The mean-field density can not change, because the equilibrium order parameter already minimises the overall repulsion energy, but the fluctuations in the density can change. The atoms can shuffle around each other, preserving their total density, but distributing themselves more regularly to remove the extra repulsion energy from pairs of atoms that are placed close to each other by chance. During this process, the excess repulsion energy will become kinetic energy.

When the atoms have shuffled in this way for some time, their repulsion energy will reach a minimum. The repulsion energy at this point might be less than in the ground state; the motion expected for a squeezed state includes times when the position fluctuations are suppressed, and the squeezed state has less potential energy than the vacuum state does. The kinetic energy will reach a maximum at these times, so the atoms won't stay where they are. The position correlations might oscillate as the excess energy transfers back and forth between repulsion energy and kinetic energy.

Figure 21 shows the squeezing ellipses for the Bogoliubov modes of a gas in a coherent state. Wave number increases from left to right. On the top line, the system size is of the same order as the healing length, and the step between wave numbers is around the reciprocal healing length. Only the lowest few modes are significantly squeezed; when the coherence function is observed over time, the discrete frequencies of these modes will be distinguishable. The bottom line shows a large system, where many modes have significant squeezing. These modes will rotate at different frequencies and get out of phase, causing the coherence function to decay and never recover.




Figure 21. Effect of gas size on the coherence oscillations.

29. Results of Wigner simulations

Section 28 made some qualitative predictions about the dynamics that would result if a gas of mutually repelling bosons was prepared in a coherent state. Some truncated Wigner simulations were carried out to test these predictions in the simplest case of a uniform, untrapped Bose gas with periodic boundary conditions. To examine the convergence of the coherent and Bogoliubov states to the ground state, the simulations were carried out with both those states as initial conditions, and were repeated for three Lieb-Liniger parameters, $\gamma = 0.01, 0.03$ and 0.1 . In these conditions, the value of $g^{(2)}$ in the exact Lieb-Liniger ground state agrees with the value in the Bogoliubov state to within a few percent [69], so the Bogoliubov graphs of $g^{(2)}$ are exact to within a few percent at $s = 0$. I am not aware of any work that has calculated the kinetic energy or the momentum distribution $n(k)$ for the Lieb-Liniger ground state. The momentum distribution has been calculated for an algebraic Bethe ansatz [74], but only for values $\gamma \gg 1$ in the Tonks-Girardeau regime.

In these simulations, the classical dynamics followed the Gross-Pitaevskii equation in one dimension, Equation 19. The number of samples required for convergence was on the

order of 10^5 . The system is homogeneous, permitting many quantities to be averaged over the length of the gas. Longer gases required fewer samples, but more work to integrate each sample. Where the graphs show a pair of parallel curves for the same quantity, these indicate the range of sampling error in the Monte-Carlo method. The software used, xSPDE, estimates time discretisation errors by repeating the integration at a smaller time step, and indicates them by plotting error bars. In all of the data plotted here, these error bars were too small to be visible.

The Wigner distribution for a squeezed state is Gaussian. This makes it fairly straightforward to draw samples for a Bogoliubov state [71, 75, 49]. The sound wave annihilation operators a_i in Equation 47 are replaced by random variables α_i to give

$$\psi(x) = \psi_0 + \sum_i u_i(x)\alpha_i - v_i^*(x)\alpha_i^*.$$

In the Bogoliubov ground state, the sound wave modes are in vacuum states. As described in Section 15, the Wigner distribution for the vacuum state is a complex Gaussian distribution, with $\overline{\alpha_i} = 0$ and $|\overline{\alpha_i}|^2 = \frac{1}{2}$, where the bars denote ensemble averages over the distribution. For an untrapped gas, the mode functions u and v are the sinusoidal ones of Equations 52 and 53, with discrete wave numbers $k_i = 2\pi i/L$. The modes for a trapped gas must be found by solving Equation 45 numerically.

The one somewhat fiddly aspect of sampling a Bogoliubov ground state is setting the grid spacing, which determines the highest momentum that occurs in the simulation. The orbitals whose states are significantly squeezed have wave numbers up to $\sqrt{2}/\xi$, and the grid must have at least a few points per healing length to include them and capture the short-time dynamics of the gas. On the other hand, the truncated Wigner method is only reliable when there is at least one particle per grid point. If this condition is not met, the sampled averages will converge very slowly. More importantly, the Moyal brackets that determine the quantum dynamics will not be well approximated by the Poisson brackets that determine the truncated Wigner dynamics. It is possible to satisfy both requirements only when there are several particles per healing length, in the coherent regime of small γ . The result that a phase space method is most reliable near the coherent limit is not especially surprising.

Some related calculations have been carried out recently [76]. The system treated was a Bose gas in a cigar-shaped harmonic trap, with parameters chosen to match the Vienna trap before splitting. The simulations are for an initial Lieb-Liniger parameter of $\gamma = 4.5 \times 10^{-4}$, which is impulsively increased to values somewhat smaller than those in the figures below. The simulations were done at finite temperatures, using a classical method and the positive-P method.

The physical situation being considered in [76] is a sudden tightening of the trap. The atoms would initially be in the ground state with $\gamma = 4.5 \times 10^{-4}$. The healing length is proportional to the transverse trap radius, according to Equation 18, so the value of γ increases as the trap tightens. Those simulations treated an increase of γ by a factor of 5,

giving a final $\gamma = 2.3 \times 10^{-3}$ and a factor of 20, giving $\gamma = 9 \times 10^{-3}$. The latter is comparable to Figure 22 below, in which $\gamma = 0.01$.

The results in [76] claim that the gas has a nonzero interaction strength initially, and this density of the gas in the harmonic trap was set to its equilibrium value for this repulsion. However, all of the calculations assumed that the state of this interacting gas was coherent. The paper acknowledges that this is only an approximation, at least implicitly, because it presents results in which nothing is supposed to have changed at the initial time, but $g^{(2)}$ equilibrates regardless. Therefore, although those results are supposed to be for a different physical process, they are in fact exactly comparable to the coherent state results in this section. The results in that paper should be treated with some suspicion, because the relaxation of $g^{(2)}$ due to the use of a coherent initial state is of the same order as the relaxation due to the change in interaction strength.

The directly comparable results occur in Figure 7 of [76], and in Figure 22 in this section. The paper has an initial $g^{(2)}(0) = 1.012$, due to a finite but small temperature. This reduces by 0.029 for a quench to $\gamma = 0.09$. The comparable result in this section is a fall in $g^{(2)}$ of 0.4 for $\gamma = 0.01$. It is plausible that this difference is due to the trapping potential in the published results.

In that paper, the correlation function $g^{(2)}(x, x')$ was computed with $x \neq x'$, and waves of correlations were seen to propagate in $x - x'$ at the Bogoliubov speed of sound. Calculating this with the truncated Wigner method would be straightforward in principle, but it would require some significant extensions to the software used in these simulations.

These simulations were implemented using the Matlab framework xSPDE [77, 78], a successor to the XMDS system [79, 80]. This framework solves a system of stochastic partial differential equations that the user specifies in a Matlab structure. It provides sensible defaults for the integration algorithm and numerical parameter, which the user can override. The input structure specifies quantities whose expectation values need to be calculated, and these quantities are automatically averaged over the ensemble and plotted with error bars.

The results are presented in the scaled units of Section 3, with the reciprocal boson density as the unit of length. This sets the scaled healing length to $\xi = 1/\sqrt{2\gamma}$. The only intrinsic time scale corresponds to the chemical potential, $s_0 = \xi^2$. Numerical values of ξ and s_0 are given in the captions.

The most basic prediction is that in the limit of small γ , both the coherent and Bogoliubov states are good approximations to the ground state. In the true ground state, all observables are steady, and the amounts by which the repulsion and kinetic energies change over time suggest how well the approximate ground state has captured the physics of the Bose gas. This is tested in Figures 22, 23 and 24. All graphs use the scaled units of Section 3, with the atomic separation as the unit of length. Axes labelled s show the time in these units. In these simulations, the gas has 200 bosons. Depending on γ , the gas has an extent of 20 to 100 healing lengths, enough to approach the infinite gas limit.

Graphs (a) show the repulsion and kinetic energies as dotted lines, and the total energy as a solid line. Each line is duplicated to indicate the sampling error. The kinetic energy T

was evaluated by taking the integral over momentum of Equation 61. The potential energy density is $V/L = \gamma(\hat{\Psi}^\dagger)^2(x)\hat{\Psi}^2(x)$, which was evaluated using Equation 59. The gas is uniform, so this is independent of x , and the plotted V/L is the average of Equation 59 over the grid points x_i . All simulations had unit condensate density; the calculated energies are slightly higher in the Bogoliubov state, which has a higher total density due to noncondensate bosons. In the coherent state, all of the bosons are condensed with zero momentum, so all of the energy is due to repulsion. As predicted, the coherent state changes over time, and some of the repulsion energy becomes kinetic energy. In contrast, the Bogoliubov state has some kinetic energy to start with, but the change over time is relatively smaller.

The total energy is conserved almost exactly. This is a feature of the truncated Wigner method [81]. The uncertainty ranges displayed in the graph are due to sampling errors, which cause the average over the samples to differ from the energy of the state from whose Wigner distribution the samples were drawn. However, the truncated Wigner dynamics exactly conserve the average over the samples.

Graphs (b) show the second order coherence function $g^{(2)}(0)$, evaluated from Equations 57, 58 and 59. This quantity is proportional to the repulsion energy, but it is normalised to the total atomic density. The value in a coherent state is 1, and values less than 1 indicate that the bosons are spread out, with pairs less likely to coincide than they would be if the bosons were placed at random. The coherent initial state relaxes quickly, and the relaxed values of $g^{(2)}(0)$ are quite close to those in the Bogoliubov state.

The time domain graphs (c) show the occupation of orbitals with nonzero momentum. The quantity plotted is the boson density $n(k)$ in momentum space, evaluated from Equation 60. Nearly all of the bosons occupy the condensate orbital with zero momentum, but $n(0)$ was set to zero when the graphs were plotted, allowing the occupation of the other orbitals to be seen.

The power spectrum of these occupations is shown in (d). If $\tilde{n}(k, \omega)$ is the Fourier transform of $n(k, s)$, under the Fourier transform of Equation 26, then $S_n(k, \omega) = |\tilde{n}(k, \omega)|^2$. The angular frequency ω is given in the same scaled units as the times s . The occupation of the condensate orbital has been set to zero so that the changes in the other orbitals can be seen on the scale of the graphs. The occupations are steadier in the Bogoliubov state than in the coherent state. This can be seen in the time graphs, and it is also apparent from the scales of the power spectra.

At small k , both power spectra are dominated by a $1/\omega$ term, which appears as a streak along the k axis. This is an artefact of the Fourier transform when the simulated range of s is shorter than the period of the sound wave mode.

Graphs (e) are power spectra of the kinetic energy density in the momentum orbitals. This is the power spectrum $S_T(k, \omega) = |k^2\tilde{n}(k, \omega)|^2$. The kinetic energy density $|k^2n(k, s)|$ was evaluated from Equation 61. This is useful as a test of convergence, because sampling errors in orbitals with large momentum can add a lot of spurious kinetic energy to the simulation. It is also a test of the Bogoliubov sound wave dynamics. The solid white line shows the Bogoliubov dispersion relation, with frequency doubled to allow for two squeezing cy-

cles per sound wave cycle. The dotted lines show the wave number $\sqrt{2}/\xi$ around which the dispersion changes from linear to quadratic. The power spectra for the coherent states fit this dispersion relation very nicely, while those for the Bogoliubov states do not. This is consistent with the prediction that the dynamics of the coherent state will be dominated by the squeezed sound wave modes, which will be vacuum states in the Bogoliubov state. The constant components at $\omega = 0$ have been set to zero to make the oscillations more visible, but their relative sizes can be seen in graphs (c).

The next question is how the discrete Bogoliubov frequencies become apparent as the gas shrinks. Figure 25 shows a gas with $\gamma = 0.03$, as in Figure 23, except that it extends over 5 healing lengths instead of 50. The results are much as expected. The sound wave mode with the longest wavelength dominates the fluctuations in energy and coherence, and its frequency is apparent in the time graphs.

Figure 26 shows a similar gas of 2.5 healing lengths. At first glance, these graphs appear to confirm the qualitative predictions too. However, graphs (c) indicate some problems with the truncated Wigner dynamics. In both initial states, the orbital with the highest momentum is empty, and at later times it is occupied by a negative number of bosons, whose negative kinetic energy cancels the kinetic energy of the other orbitals. Negative bosons are entirely unphysical in this non-relativistic system. However, the Wigner moment for the particle number includes a correction for symmetric order, which subtracts half a particle from the orbital. In the initial state, the random amplitude of the Monte-Carlo order parameter is equal to this correction, and they cancel to give an occupation of zero. At later times, the random amplitude has decayed, and subtracting the correction gives a negative occupation number. The truncated Wigner approximation is apparently not valid for these parameters, but the reasons for that are unclear.

These results are interesting in the context of elongated condensates. In a numerical treatment of these systems, the simplest way to initialise the state of the condensate is to use the equilibrium order parameter as the amplitude of a coherent state. No doubt this is safe when the trap is so tight that the atoms are restricted to the ground orbital of the trapping potential along the tight axes. However, it is often the case that the repulsion energy is of the same order as the zero point energy, so that the equilibrium order parameter is spread out somewhat further than the ground orbital. The atoms are still tightly trapped, and the extent of the condensate along the tight axis is short, so this is a situation where the oscillations in coherence could be quite large. The Bogoliubov state for the atoms along the tight axis will differ significantly from the coherent state. This is suggested by some simulations that have been done with the potential for the Vienna trap.

These simulations can be compared to some quantities calculated for the exact Lieb-Liniger ground state. The momentum density $n(k)$ has been computed by an ABACUS method [74]. The results of that calculation are compared to the Bogoliubov density of Equation ? in Figure 27. The published results are for $\gamma \geq 0.25$, so they are not directly comparable to these simulations. However, the Bogoliubov approximation should improve as γ decreases.

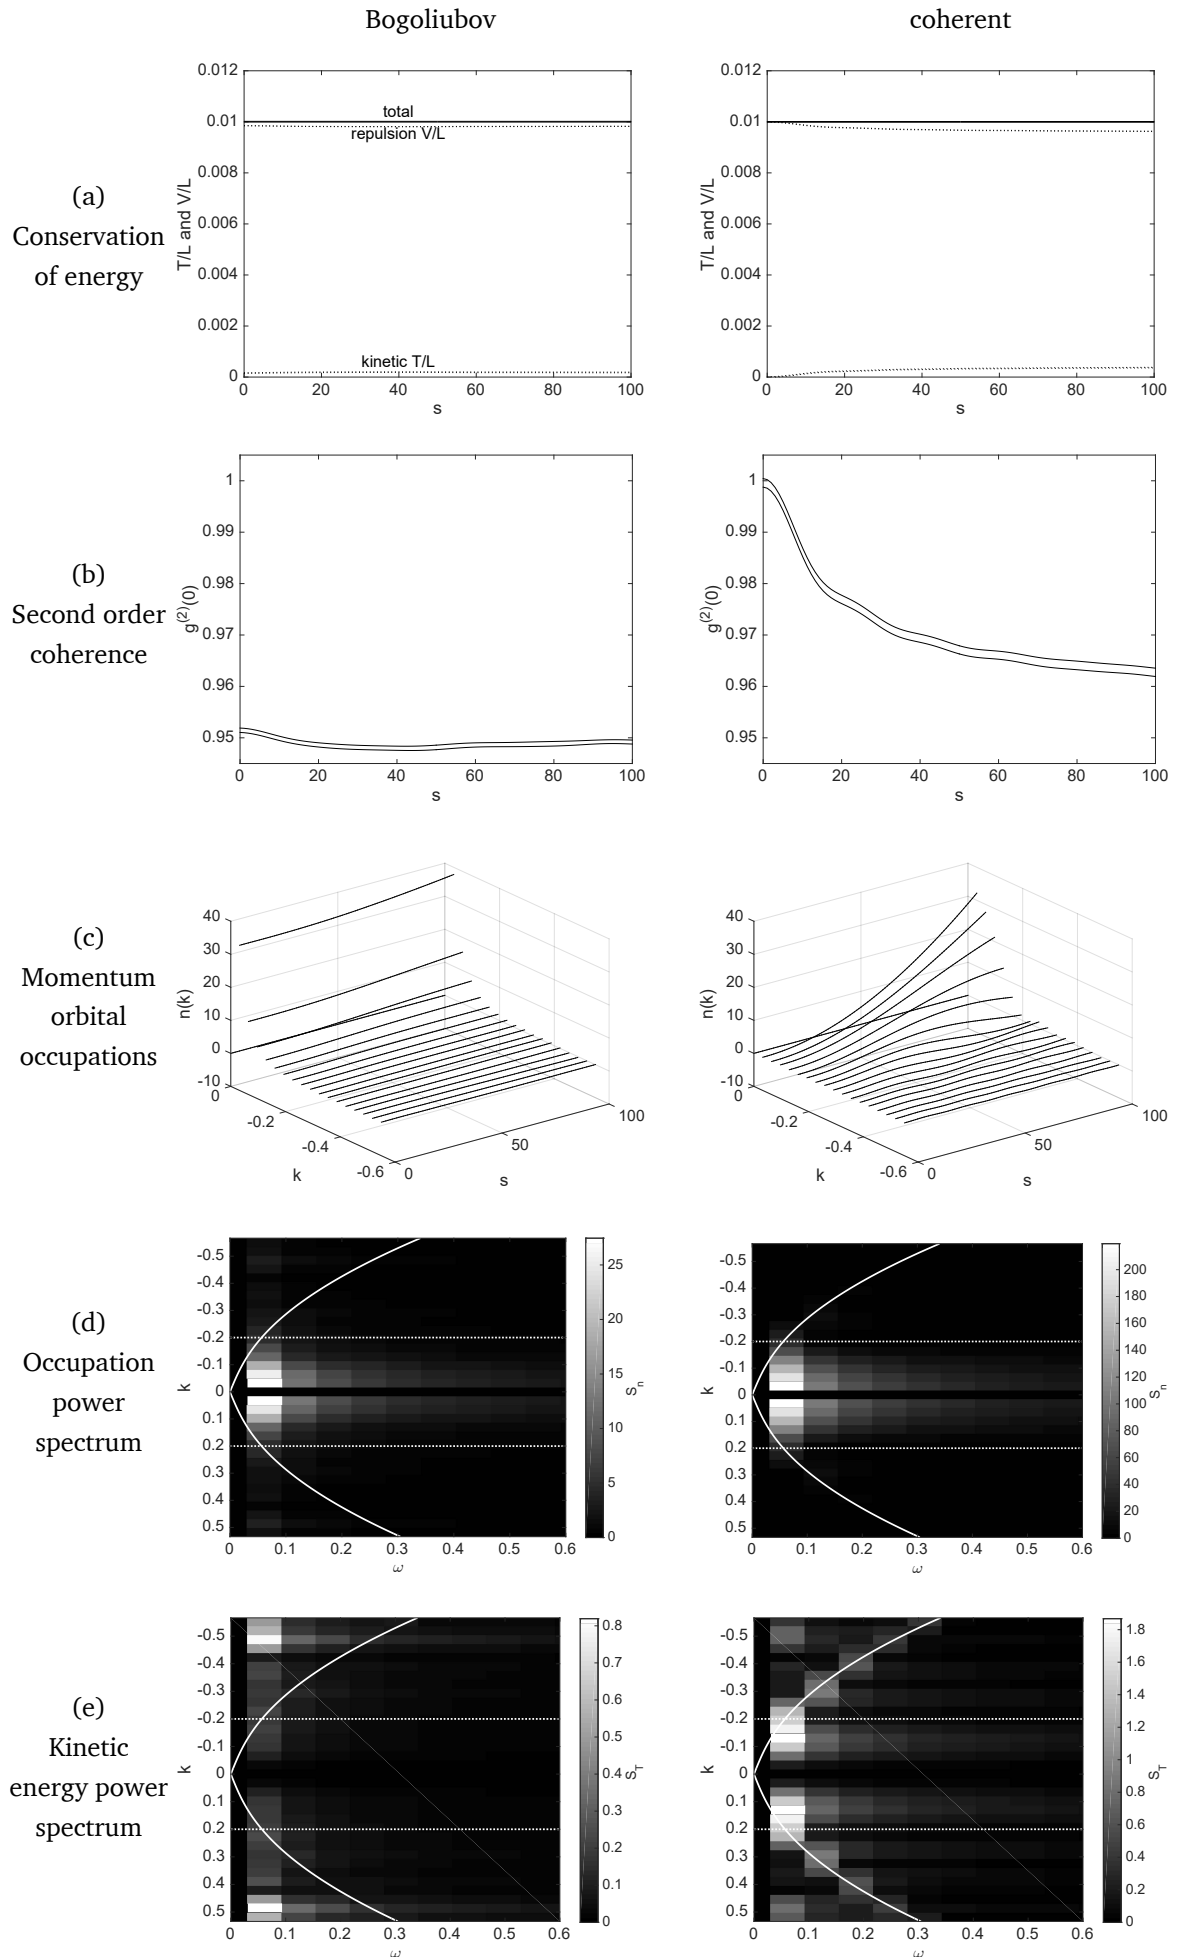


Figure 22. Comparison of initial states, $\gamma = 0.01$, length 200, $\xi = 7.1$, $\xi^2 = 50$.

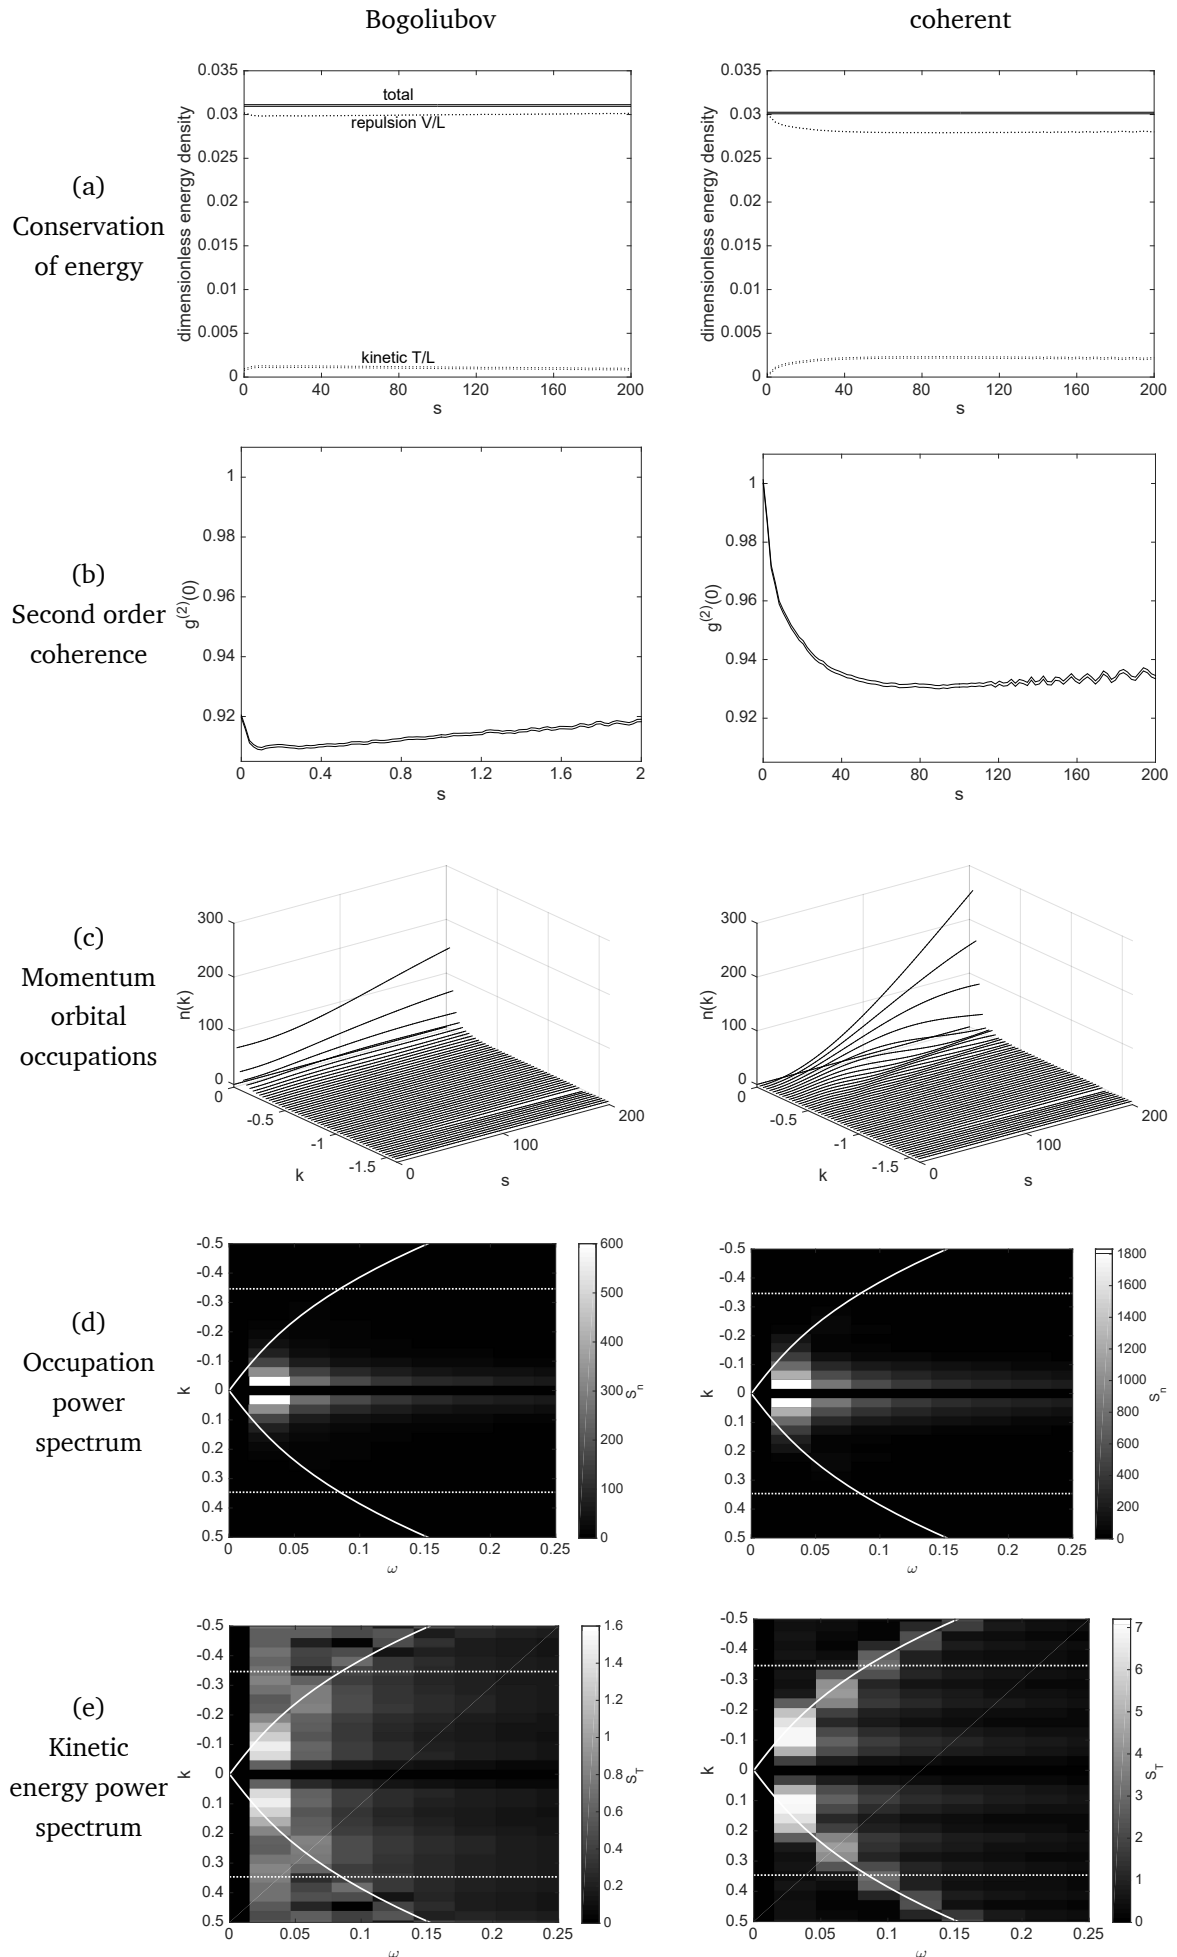


Figure 23. Comparison of initial states, $\gamma = 0.03$, length 200, $\xi = 4.1$, $\xi^2 = 14$.

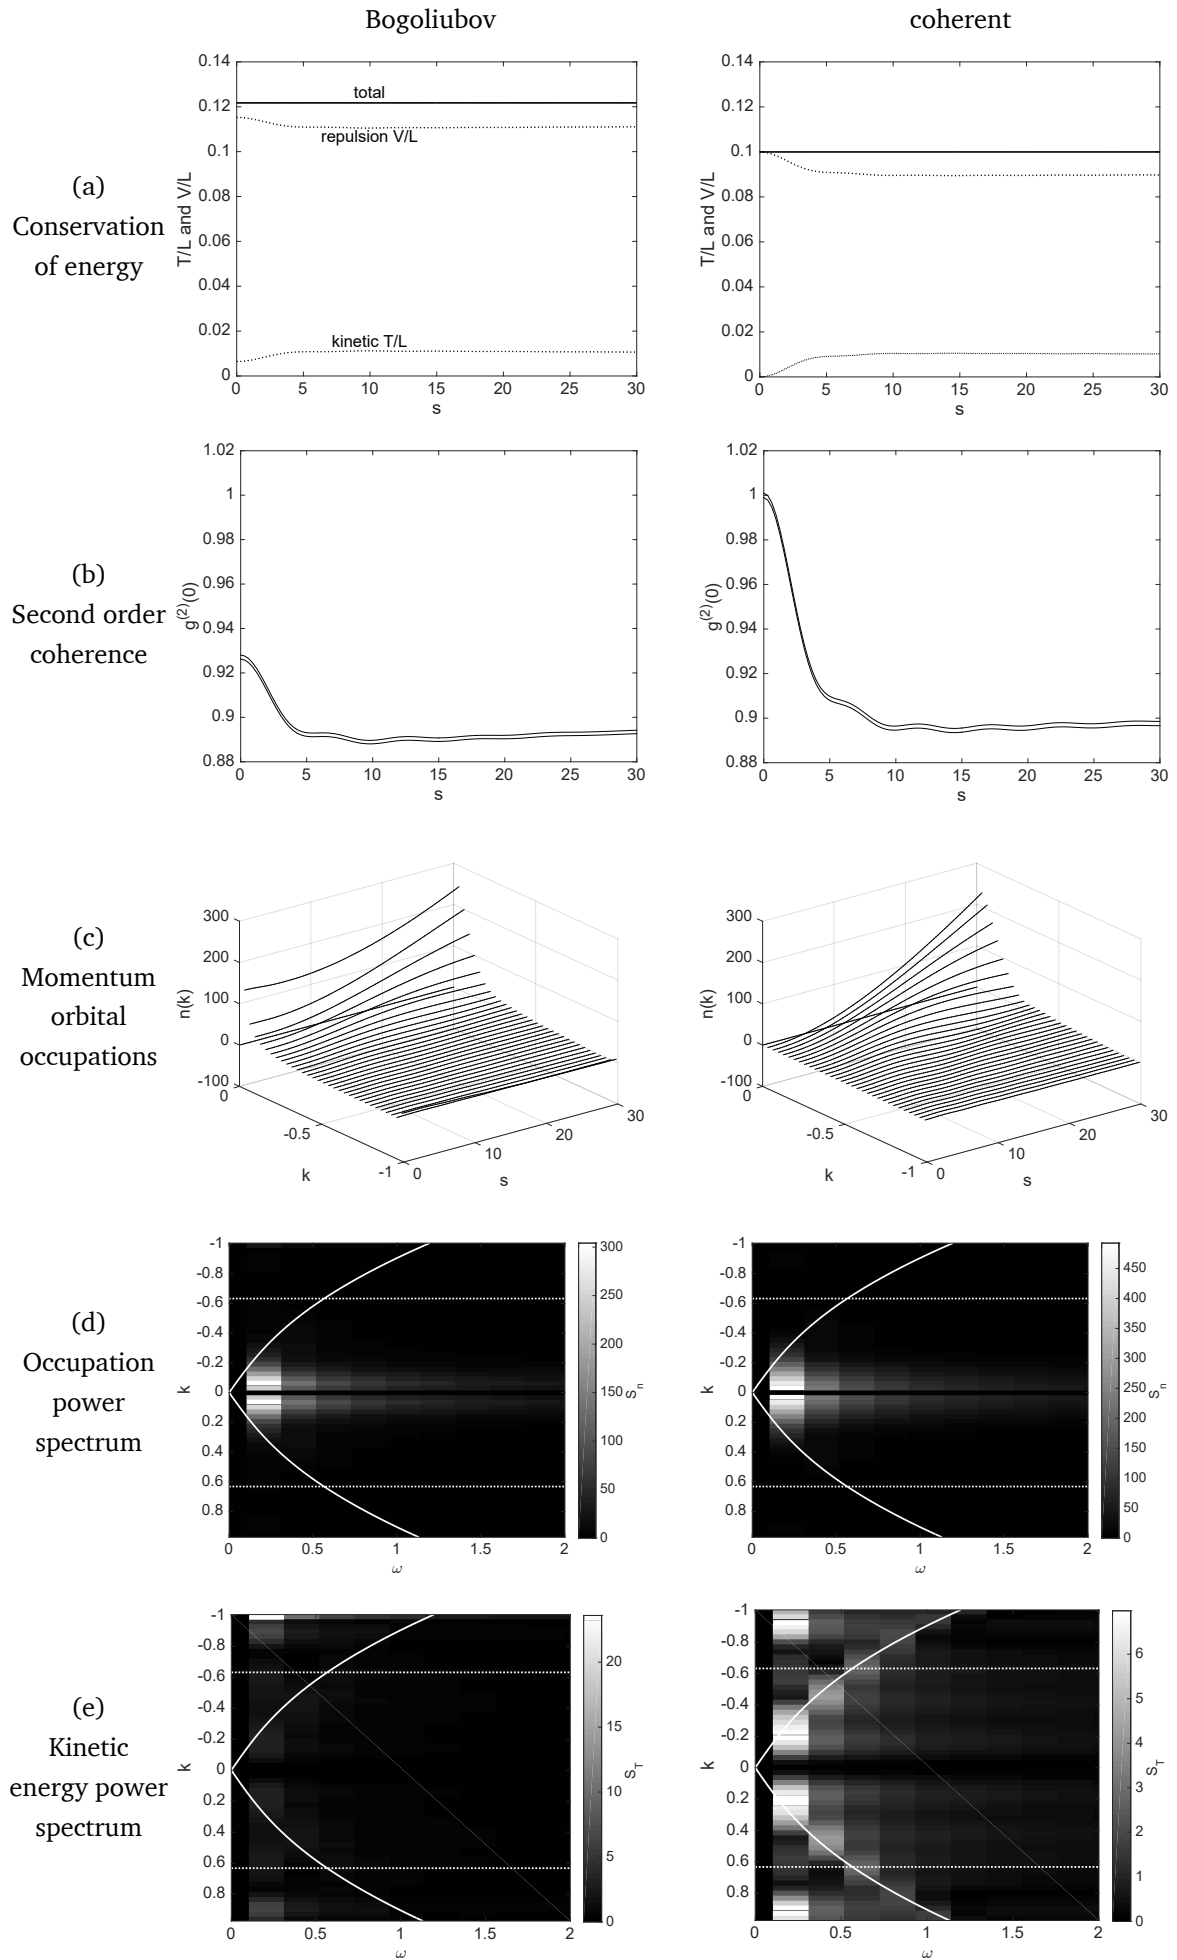


Figure 24. Comparison of initial states, $\gamma = 0.1$, length 200, $\xi = 2.2$ and $\xi^2 = 5$.

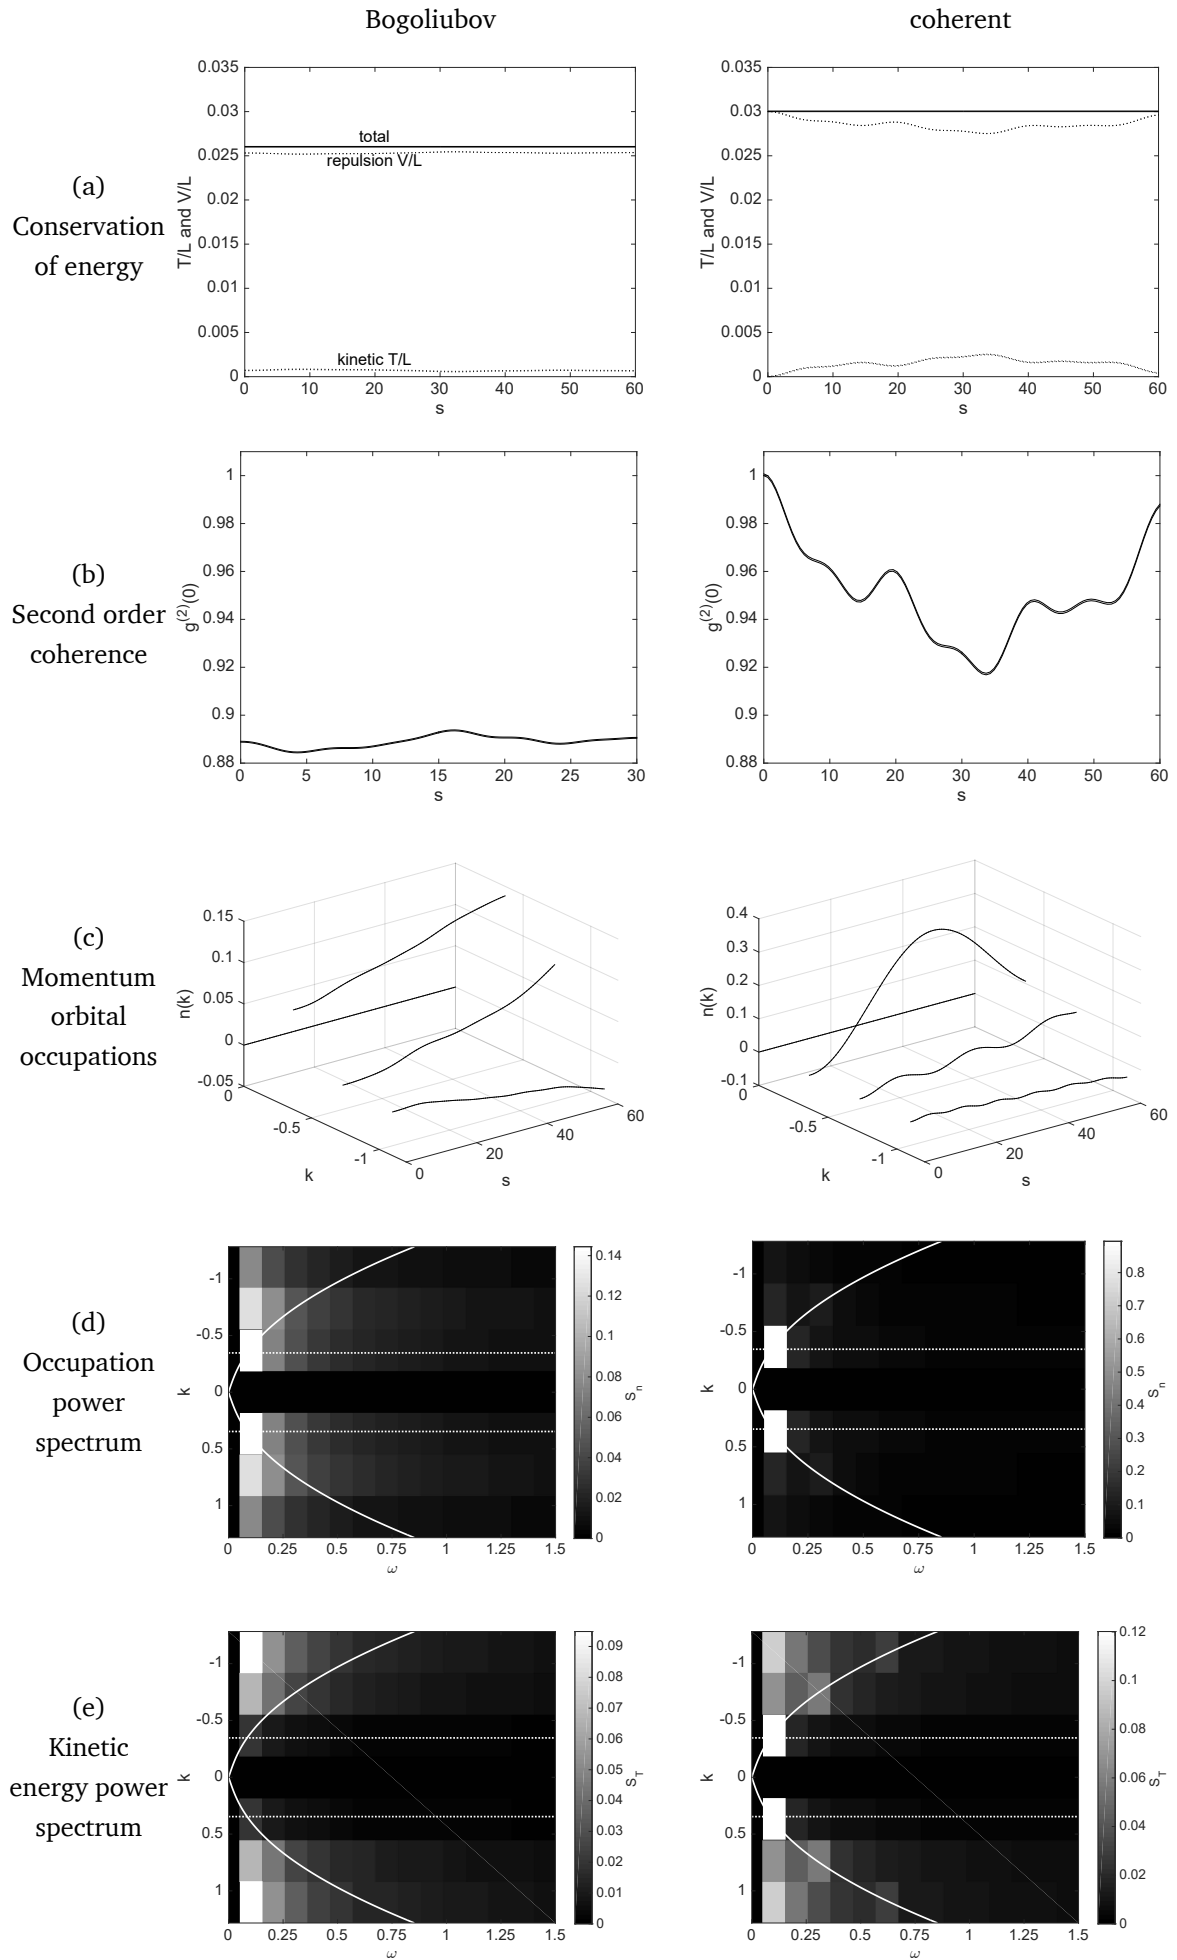


Figure 25. Oscillating coherence with $\gamma = 0.03$, length 20, $\xi = 4.1$, $\xi^2 = 14$.

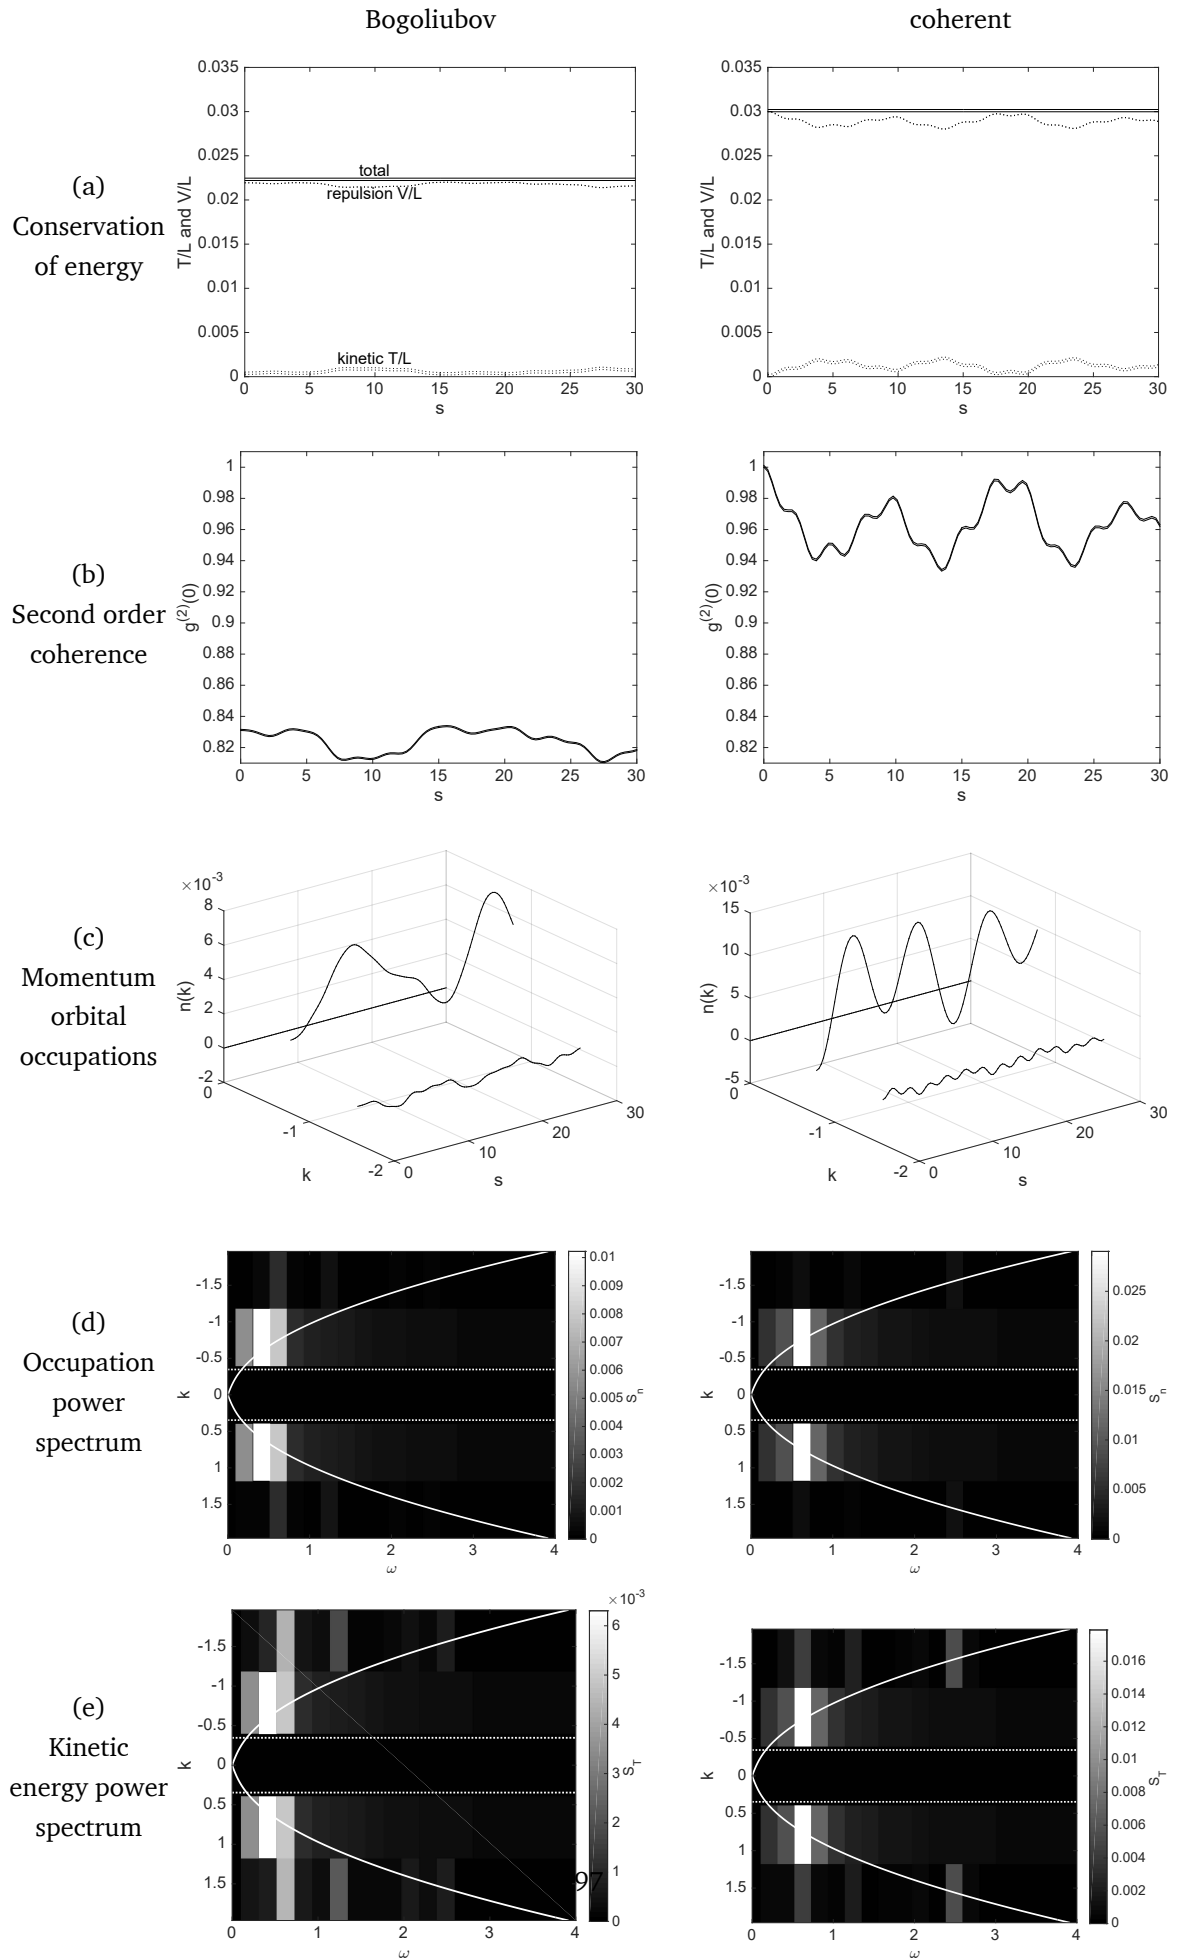


Figure 26. Oscillating coherence with $\gamma = 0.03$, length 10, $\xi = 4.1$, $\xi^2 = 14$.




Figure 27. Momentum density by Bogoliubov (solid line) and ABACUS (crosses) for $\gamma = \frac{1}{4}$ (bottom line), $\frac{1}{2}$ and 1 (top line).

The two methods give similar densities for $\gamma = 0.25$. The main difference is that the exact values are consistently smaller than the Bogoliubov ones. This could be explained by the non-condensate density in the Bogoliubov approximation. The exact results are for a gas with unit total density, but the Bogoliubov gas has unit condensate density, so its total density is somewhat larger than 1. The Bogoliubov gas is more dense, which would explain why its momentum density is higher.

30. Correlations between numbers oscillating in a trap

The motivation to investigate the dynamics in Section 29 came from attempts to simulate the splitting of the Vienna trap, with a coherent initial state. These simulations have not been completed, since it became apparent that the correlations in atom positions were more complicated than they initially appeared. This section presents some of the initial results. The rest of the chapter presents some formalism developed for coherent states. This will need to be revised to allow the use of Bogoliubov states or other more accurate approximate ground states before it can be used for numerical work.

Figure 28 shows a simulation of the balance of particle numbers between the two sides of the split trap. The observable plotted is related to the number difference observables of Section 27. It has not been placed in normal order, rather it has been normalised by the total number of atoms. This means that the expectation value when the two split components are independent coherent states is 1 instead of 0.

The simulation was carried out by a truncated Wigner method. The trap was treated in two dimensions, with the quartic fit potential taken in the x - z plane. This is the poten-

tial shown in Figure 19. The classical equation of motion is Equation 17, with $l_0 = x_0$, the characteristic oscillator length along the x and y axes of the initial trap. The equilibrium order parameter was found by solving the Gross-Pitaevskii equation in imaginary time with the initial trap potential, then Wigner samples were taken for a coherent state with this amplitude and the dynamics computed. The error bars show the statistical error from sampling the Wigner distribution. The axis labelled t shows the scaled time, in units of μm^2 . This scaled time was labelled s in the one-dimensional simulations shown in Section 29. The numerator of the plotted observable was computed from a moment similar to Equation 56, with some ordering corrections. The denominator was computed from the moment in Equation 58.

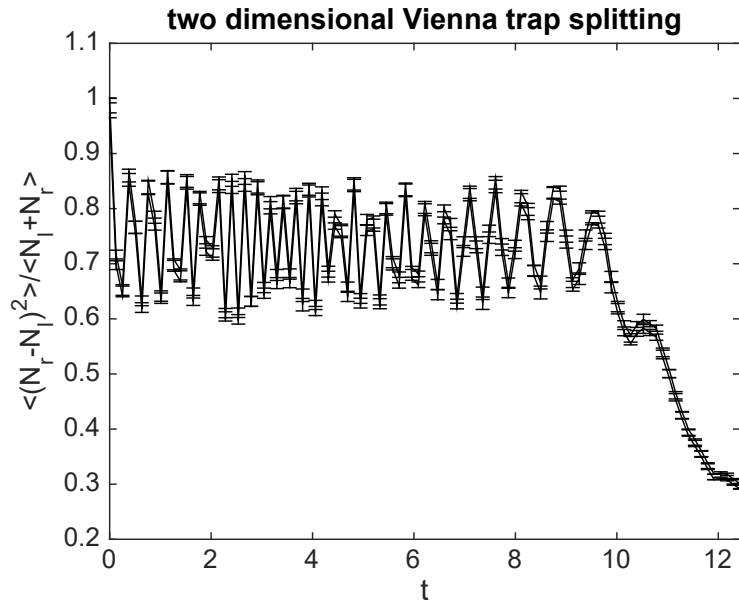


Figure 28. Oscillations in numbers of atoms on either side of a splitting trap.

Initially, the atoms are assumed to have condensed; the atomic field is a coherent state, in the equilibrium order parameter of the cigar-shaped trap. The y - z plane divides the atomic cloud into two equal parts. In the initial state of the atoms, the states of both parts are independent coherent states, and the number difference observable has expectation value 1. This is the case where the atoms are distributed at random, and the numbers on either side of the cloud are binomially distributed.

This correlation observable immediately starts to oscillate, indicating that the atoms do not remain in a coherent state. At first, the numbers of atoms on each side of the trap become more nearly equal. Then the excess repulsion energy from the coherent state compared to the atomic ground state converts back and forth between position correlations and, presumably, kinetic energy. The frequency of the oscillations reduces as the trap starts to split; the first part of the splitting involves the trap elongating along the splitting axis, which will cause the trap frequency to reduce.

As the two sides of the trap start to separate, and the trap enters the tunnelling regime, the number difference decreases rapidly. The cloud splits into two parts around $t = 11$ to 12, corresponding to the $15 \mu\text{s}$ and $16 \mu\text{s}$ columns in Figure 19. The number difference takes a steady value in the final state. Once the two clouds have entirely separated, no atoms can pass from one component to the other, and the oscillations must stop.

If the atoms did not interact, the atomic field would remain in a coherent state during the whole splitting process. This is a tensor product of coherent states for the fields on either side of the trap, so the number difference observable would remain constant. In the simulation, the atoms repel each other. Before the cloud splits, it spreads out along the splitting axis. This provides room for the positions of the atoms to become anticorrelated, so that they avoid colliding with each other and have a reduced interaction energy. The anticorrelated positions mean that, if one atom is on the left side of the trap, other atoms are likely to be on the right side of the trap. This reduces the difference between the numbers on the two sides.

The frequency of the oscillations in the simulation is about $3 \mu\text{m}^{-2}$, while the trap frequency, $1/x_0^2$ is around $30 \mu\text{m}^{-2}$. It is not obvious why the oscillations in the position and number correlations are so much slower than the motion of an atom.

These oscillations are relevant to the Luttinger liquid model for the split trap. This model is initialised by assuming that the transverse state the split gas is coherent, with both components having equal amplitudes. If this were the case, a coherent state would be a steady state of the trap, at least along the splitting axis. On the other hand, the number observable in this simulation oscillates with a large amplitude. This suggests that a coherent state is not a very close approximation to the actual transverse state, and that the suppressed number fluctuations due to repulsion are a plausible explanation for the excess phase fluctuations in the experiment compared to the Luttinger liquid model.

31. Wigner moments for interference observables

When the atoms are released from the split trap, the density quickly reduces to a level where the repulsion energy is small compared to the kinetic energy, and after this each component expands freely; by the time the atoms are photographed, the clouds overlap, and in the transverse plane they extend further than both of the trapped components did. On the other hand, the atoms are trapped loosely along the z -axis, so they have lower wave numbers, and expand slowly. Suppose the interval from when the atoms are released until their photograph is taken is s in scaled units. The transverse velocity of the atoms, in scaled units, is their wave number k . So the density in the photographed cloud at position (x, y, z) is about $|\tilde{\psi}(x/s, y/s, z)|^2/s^2$, where the order parameter $\psi(x, y, z)$ has been expanded over longitudinal position and transverse momentum as the Fourier transform

$$\psi(x, y, z) = \int \langle x, y, z | k_x, k_y, z \rangle \tilde{\psi}(k_x, k_y, z) dS = \frac{1}{2\pi} \int e^{i(k_x x + k_y y)} \tilde{\psi}(k_x, k_y, z) dS.$$

The discussion of the split trap will involve one, two and three-dimensional integrals over the gas cloud. These will be abbreviated with $dS = dx dy$, and $dV = dx dy dz$.

The photograph is taken along the y axis, so it records a column density

$$\frac{1}{s^2} \int |\tilde{\psi}(x/s, y/s, z)|^2 dy = \frac{1}{s} \int |\Psi(x/s, y, z)|^2 dy.$$

The integral has been moved from k_y to y using Rayleigh's theorem, and the Fourier transform is

$$\psi(x, y, z) = \int \langle x, y, z | k, y, z \rangle \Psi(k, y, z) dk = \frac{1}{\sqrt{2\pi}} \int e^{ikx} \Psi(k, y, z) dk.$$

In the quantum theory, the observable measured when the interference fringes are photographed is the atomic density at position z with transverse momentum k ,

$$\rho(z, k) = \int \hat{\Psi}^\dagger(k, y, z) \hat{\Psi}(k, y, z) dy.$$

In the experiments [28], this density produced interference fringes along the k axis. The phase of these fringes varied with z . In each photograph, a phase function $\phi(z)$ was fit to the fringes by least squares, then the repeated photographs were averaged to give a phase correlation function $C(z, z', t) = \text{Re}\langle e^{i\phi(z)} - i\phi(z') \rangle$.

In order to simulate these measurements, this phase correlation function must be related to an observable of the atomic field. For simplicity, suppose that the envelope of the interference fringes is a known function $f(k)$. Otherwise, this envelope would have to be determined by non-linear optimisation, and the observables corresponding to the fitted parameters would be very complicated. Also suppose that the fringe spacing is known, and the fringe density is a sinusoidal function of bk , where the length b is the separation of the parallel lines in the split trap. This was $3 \mu\text{m}$ in the experiment.

Under these assumptions, the density can be written as a linear combination of three functions of k , whose coefficients are functions of z .

$$\rho(z, k) = f(k) \left(\frac{u_1(z)}{\sqrt{2}} + u_2(z) \sin bk + u_3(z) \cos bk \right).$$

The contrast and the phase can be determined straightforwardly from the vector $u(z)$; the aim is to find observables for the values of these coefficients that fit the photographed density in a least squares sense. The two components in the split trap are confined tightly along the splitting axis, so the envelope $f(k)$ will be broad compared to the fringe spacing b and the functions $f(k)/\sqrt{2}$, $f(k) \sin bk$ and $f(k) \cos bk$ will be nearly orthogonal. If f is normalised appropriately, the least squares solution will be

$$u_1(z) = \int \rho(z, k) dk \quad u_2(z) = \int \rho(z, k) \sin bk dk \quad u_3(z) = \int \rho(z, k) \cos bk dk.$$

The phase can now be determined as

$$e^{i\phi(z)} = u_2(z) + iu_3(z) = \int \rho(z, k) e^{ibk} dk.$$

The correlation function is

$$C(z, z', t) = \text{Re}\langle e^{i\phi(z) - i\phi(z')} \rangle = \text{Re}\left\langle \int \rho(z, k) e^{ibk} dk \int \rho(z', k') e^{-ibk'} dk' \right\rangle.$$

Although this observable is fourth-order in the creation and annihilation operators, finding its Wigner function is straightforward. From the convolution theorem,

$$\int \rho(z, k) e^{-ibk} dk = \int \hat{\Psi}^\dagger(z, k) \hat{\Psi}(z, k) e^{-ibk} dk = \sqrt{2\pi} \int \hat{\Psi}^\dagger(z, x) \hat{\Psi}(z, x + b) dx.$$

For $z \neq z'$, the fields inside the expectation value commute, so this moment is already symmetric, and its Wigner function is

$$C(z, z', t) = \text{Re}\left\langle \int \psi^\dagger(z, x) \psi(z, x + b) dx \int \psi^\dagger(z', x') \psi(z', x' + b) dx' \right\rangle.$$

So the obvious method, computing a correlation function for the relative phase of the Wigner samples, will give the right answer except at $z = z'$.

32. Hartree modes

The discussion of the split trap experiments started in Section 24 with a line of cold atoms that split into two lines. By Section 31, the description had changed to a two-dimensional plane. This section adapts that theory to more closely resemble the original description. There are several aims to this. It would be nice to know if the original description in terms of lines of gas and beam splitting is quantitatively true. When the experiments are analysed numerically, it will be more efficient to describe them in one dimension than in two. Furthermore, the truncated Wigner method is most trustworthy when many atoms occupy few modes; a model that employs a handful of discrete transverse modes will be more reliable than one on a transverse grid.

The remainder of this chapter derives one way to expand the order parameter over a small number of transverse modes, to reduce the Gross-Pitaevskii equation to a system of one-dimensional partial differential equations for the components of those modes, and to initialise those components so that the expanded three-dimensional field is distributed according to the Wigner distribution for a coherent atomic field. The derivation is fairly complete, but no simulations have yet been done using this approach. The results in Section 30 suggest that a coherent state is not an accurate enough approximation to the ground state of the interacting gas before the trap is split, and it is likely that this work would need to be extended to a Bogoliubov initial state to avoid the oscillating correlations in those results.

Figure 19 shows the transverse ground state of the atomic field as the trap splits. The transverse potential becomes a double well; each well has its own ground state, and these are nearly orthogonal. The initial ground state becomes a symmetric superposition, and the vacuum mode becomes an antisymmetric superposition. The energy gap between the ground state and the vacuum mode reduces as the well expands and splits; eventually, the

symmetric and antisymmetric states become degenerate. The virtual atoms in the vacuum mode become a phase difference between the two lines of condensate. If this field is to be expanded over a small number of transverse modes, it is clear that the modes will have to evolve with the trap potential.

The next step is to find mode functions that satisfy this description. The quartic potential that approximates the trap has the form $K^2(x, y, z) = Q^2(z) + R^2(x, y)$. This allows a complete set of transverse modes θ_j to be defined, as the eigenmodes of Schrödinger's equation in two dimensions,

$$(-\partial_{xx} - \partial_{yy} + R^2) \theta_j = k_j^2 \theta_j. \quad (62)$$

The wave function can be expanded in the form

$$\psi(s, x, y, z) = \sum_j \psi_j(s, z) \theta_j(x, y). \quad (63)$$

When this form is substituted into the time dependent Schrödinger's equation, it reduces to a set of independent one-dimensional differential equations,

$$2i\partial_s \psi_j = (-\partial_{zz} + Q + k_j^2) \psi_j.$$

The thermal energy available in the bisection experiment is only sufficient to excite atoms to the lowest few transverse modes, so only a few of these equations would have to be solved, and all the other ψ_j remain zero.

This works for non-interacting atoms, but the order parameter for interacting bosons satisfies a nonlinear Gross-Pitaevskii equation instead of the linear Schrödinger's equation. The input and output modes are still well defined, being the stationary states with lowest energy of the initial and split traps; these can be found by standard numerical methods. On the other hand, the vacuum mode becomes a vague concept when the atoms interact. This is supposed to be an excited state of the initial trap, but it is not even clear that excited stationary order parameters exist when atomic repulsion is allowed for. Even if the vacuum mode can be defined, the solutions to the GPE do not obey a superposition principle, so the input and vacuum modes might not remain independent as they propagate through splitting of the the trap.

It is physically clear that some one-dimensional model ought to exist, particularly in the limit where all the atoms are in the transverse ground state; the conceptual and computational benefits of treating the atoms in one dimension are significant. This section derives a set of modes that are self consistent with respect to atomic repulsion, similar to the Hartree states for mutually repelling electrons in an atom. Section 33 will expand the dynamics of the order parameter over these modes. The resulting equations of motion are not exactly independent, but the interactions between modes are reduced as much as possible.

The first step in defining the self-consistent modes is to allow the order parameter to separate approximately instead of exactly, and the modes to vary over time. Both of these can be allowed by expanding the order parameter in the form

$$\psi(s, x, y, z) = \sum_j \psi_j(s, z) \theta_j(s, x, y, z). \quad (64)$$

The new expansion reduces to Equation 63 in the case that ψ_j and θ_j are independent of time, while θ_j is independent of z . We will relax these requirements, requiring the functions to vary gradually with these variables instead of being independent of them, so that the derivatives $\partial_s \psi_j$, $\partial_s \theta_j$ and $\partial_z \theta_j$ are small but not necessarily zero. When the dynamical equations are derived in the next section, the rates at which atoms scatter between the modes will turn out to be proportional to these derivatives.

The more general form of Equation 64 affords some extra freedom in the choice of modes: an arbitrary function of z can be factored from ψ_j into θ_j without changing ψ . These functions can be normalised by imposing a condition

$$\int \theta_i^* \theta_j dS = \delta_{ij}. \quad (65)$$

With this normalisation, the total number of atoms is

$$\int |\psi|^2 dV = \sum_j \int |\psi_j|^2 dz.$$

In order for $\partial_z \theta_j$ to be small, the phase of θ_j needs to vary gradually along the trap axis; apart from that, this phase will remain arbitrary. With various levels of numerical effort, it could be set to zero on the trap axis, chosen to minimise the kinetic energy of the wave function θ_j , or chosen to minimise the rates at which atoms scatter between modes.

In the next section, it will be shown that the wave functions ψ_j evolve most independently when the modes θ_j satisfy a transverse Schrödinger equation

$$(-\partial_{xx} - \partial_{yy} + K^2 + 8\pi a\rho)\theta_j = (k'_j)^2 \theta_j, \quad (66)$$

where $\rho(s, x, y, z)$ is a real valued function representing the atomic density, and the modes satisfy a self-consistency condition

$$\rho = |\psi|^2 = \left| \sum_j \psi_j \theta_j \right|^2. \quad (67)$$

This condition is very similar to the one Hartree used to approximate the wave function of an atom whose electrons repel each other. That is not entirely surprising: Hartree's simplifying assumption, that every particle occupies a well-defined state, is also assumed in deriving the Gross-Pitaevskii equation. It turns out that his iterative method of finding stationary states of an atom is also able to find excited stationary states of a condensate order parameter. The results were shown above in Figure 19.

Section 22 noted that the one-dimensional healing length in the trap, $\xi = 0.29 \mu\text{m}$, is comparable to the transverse length scale $x_0 = 0.24 \mu\text{m}$. Therefore repulsion will significantly expand the cloud in the transverse plane.

The mean-field theory for a condensate has the additional constraint that the state of every atom is the same, and in this way our problem is simpler than Hartree's. However, in the usual atomic problem, only the ground state must be determined. In the bisected trap,

at least one of the excited states is also involved, the one that mixes vacuum noise into the antisymmetric mode. These two modes must be found self-consistently, as a function of the spin variable that represents the distribution of atoms between them. For a Monte-Carlo simulation, the self-consistent modes can be computed once, but the longitudinal dynamics are simulated many times.

Hartree started by constructing a wave function in which N electrons occupied the lowest $N/2$ orbitals of a hydrogen atom [82]. This is not a stationary state when electrostatic repulsion is taken into account, so he had to adjust it. First, he added up the density of charge in these orbitals, along with the nuclear charge, to find a total charge distribution and the corresponding electrostatic potential. He took this potential as given, and found adjusted orbitals for a single electron moving under its influence. Finally, he constructed a new wave function, with electrons in the lowest $N/2$ of the adjusted orbitals, and repeated the process. This converged to a many-electron wave function, self-consistent under the assumption that the state separated with each electron occupying a definite orbital.

The trapped bosons in the split trap will be distributed between the two transverse modes with lowest energy, leaving the degrees of freedom as the linear density in each mode, and the relative phase between them. The other modes have significantly higher energies, but the lowest two become degenerate as the trap splits. For fixed densities and phase, it is possible to search for two mode functions that produce a self-consistent transverse density when combined with the appropriate phase. The shapes of these modes depend not only on the proportions in which atoms are distributed between them, but also on the total linear density of atoms at that point along the trap. However, the atoms repel strongly, so the total linear density will stay close to its equilibrium value, and a linear expansion around equilibrium should suffice.

These modes can be computed iteratively. Fortunately, the iterations converge, and the transverse Gross-Pitaevskii equation has excited stationary states.

33. One-dimensional field dynamics

Section 32 derived sets of self-consistent transverse modes for the interacting atomic field. These modes depend on a one-dimensional order parameter that has multiple components, one for each transverse mode. This section derives equations of motion for these components, reducing the Gross-Pitaevskii equation to a system of partial differential equations in one space coordinate. The effective healing length in these equations varies along the axis, depending on the linear density of the atoms and their distribution among the transverse modes. In the limit where the transverse modes are uniform along the line, and the linear density is constant, these equations become independent. Near this limit, atoms tunnel at small rates between the field components.

Suppose some atoms are contained by an external potential $K^2(x, y, z)$, and their order parameter satisfies Equation 15,

$$2i\partial_s\psi = (-\nabla^2 + K^2 + 8\pi a|\psi|^2)\psi.$$

The potential K^2 has a sharp minimum around the z axis, where the particles are trapped, and it varies gradually along the length of the trap. Let $\rho(s, x, y, z)$ be a notional density of atoms, define modes $\theta_j(s, x, y, z)$ that satisfy Equation 66, then expand ψ in the form of Equation 64. At this point, ρ is an arbitrary function. Following Hartree, we will eventually specify that the atoms occupy a fixed superposition of the transverse modes, and require ρ and the mode functions θ_j to be mutually consistent according to Equation 67. Their orthogonality and normalisation were discussed in Section 32.

When the expansion of Equation 64 is substituted into Equation 15, and terms that occur in Equation 66 are collected, the result is

$$2i \sum_j ((\partial_s \psi_j) \theta_j + \psi_j (\partial_s \theta_j)) \\ = \sum_j \left(-(\partial_{zz} \psi_j) \theta_j - 2(\partial_z \psi_j) (\partial_z \theta_j) - (\partial_{zz} \psi_j) \theta_j \right. \\ \left. + \psi_j (-\partial_{xx} - \partial_{yy} + K^2) \theta_j + 8\pi a |\psi|^2 \psi_j \theta_j \right).$$

The aim is to eliminate the the transverse coordinates x and y , leaving a system of partial differential equations in s and z , that will hopefully resemble the Gross-Pitaevskii equation.

Projecting out the transverse mode θ_i on each side, and replacing the transverse derivatives with their eigenvalue from Equation 66, gives

$$2i(\partial_s \psi_i) = (-\partial_{zz} + k_i^{2'}) \psi_i + 8\pi a \sum_j \psi_j \int \theta_i^* \theta_j (|\psi|^2 - \rho) dS \\ - \sum_j \left(\psi_j \int \theta_i^* (\partial_{zz} + 2i\partial_s) \theta_j dS + 2(\partial_z \psi_j) \int \theta_i^* (\partial_z \theta_j) dS \right). \quad (68)$$

The longitudinal wave functions ψ_j are independent of the transverse coordinates, so they factor out of the integrals. The second line is small compared to the first line, because every term is proportional to a derivative of θ_j with respect to z or s .

The self-consistency condition $\rho = |\psi|^2$ eliminates the term in Equation 68 that is proportional to $8\pi a$, causing the order parameters ψ_j to evolve as independently of each other as possible. Given this condition, the eigenvalue $k_j^{2'}$ is the expectation value

$$k_j^{2'} = \int \theta_j^* (-\partial_{xx} - \partial_{yy} + K^2 + 8\pi a |\psi|^2) \theta_j dS.$$

Substituting this eigenvalue, and neglecting the second line of Equation 68 for now, gives

$$2i\partial_s \psi_j = \left(-\partial_{zz} + k_j^2 + 8\pi a \int |\theta_j|^2 |\psi|^2 dS \right) \psi_j, \quad (69)$$

where

$$k_i^2 = \int \theta_j^* (-\partial_{xx} - \partial_{yy} + K^2) \theta_j dS.$$

This is the transverse energy of θ_i , an eigenstate of the Hartree potential $K^2 + 8\pi a \rho^2$, in the trap potential K^2 .

What have we gained, in return for the extra work of computing transverse modes? Equation 69 is still three-dimensional, with an integral over the $|\psi|^2$ that we started with, so the modes are not entirely separated. On the other hand, in the absence of the second line of Equation 68, the atoms that start out in each mode remain in that mode: they collide with atoms in other modes, and those collisions push them along the z axis, but they do not scatter into other modes.

This is simplest in the truly one-dimensional case, where the energies k_j^2 for the excited states are too large for those states to matter, so that $\psi = \psi_0\theta_0$. It is shown below that, provided these energies are large enough and the trap potential changes adiabatically, the second line of Equation 68 has a small effect. Equation 69 preserves the form $\psi = \psi_0\theta_0$, so when this effect is neglected, the particles remain in the original mode. The result is a one-dimensional Gross-Pitaevskii equation with a z -dependent repulsion,

$$2i\partial_s\psi_0 = \left(-\partial_{zz} + k_0^2 + \frac{4a}{l_0^2(z)}|\psi_0|^2 \right) \psi_0. \quad (70)$$

As in Chapter 2, the nominal trap radius is $l_0^2(z) = (2\pi \int |\theta_0|^4 dS)^{-1}$, and the ground orbital is now self-consistent with respect to repulsion. The integral $\int |\theta_0|^4 dS$ depends on ψ_0 through the self-consistency condition; this integral depends on the two real variables z and $|\psi_0|^2$, and could easily be tabulated. The linear density of particles $|\psi_0|^2$ will often vary only by small amounts, in which case a linear expansion about its mean value would suffice, and the result would be a Gross-Pitaevskii equation with an extra term in $|\psi_0|^4\psi_0$.

In the split trap experiments, the order parameter can be represented in this single-mode form before it starts to split. After the two wells have split entirely, the lowest two transverse modes become degenerate, so the symmetric and antisymmetric modes can be replaced by the ground states θ_l and θ_r of the two wells. Since these do not overlap, Equation 69 reduces to two independent equations of the form 70 for the corresponding fields ψ_l and ψ_r .

The interesting part of the trap splitting is the tunnelling regime. The transverse tunnelling rate that defines this regime is equivalent to the energy gap between the ground and vacuum modes. In the tunnelling regime, the gap is too small for the changes in the trap potential to be adiabatic, and the vacuum mode matters. On the other hand, tunnelling matters too: the modes are not degenerate, so they can not be converted into ground states of two separate wells. In this case, the system must be represented using the input and vacuum modes θ_0 and θ_1 . Equation 69 becomes a pair of coupled equations for ψ_0 and ψ_1 , with nonlinear terms of the form

$$2i\partial_s\psi_0 = \dots + 8\pi a \left(\sum_{ij} \psi_i^* \psi_j \int |\theta_0|^2 \theta_i^* \theta_j dS \right) \psi_0,$$

and likewise for ψ_1 . The integrals are functions of ψ_0 , ψ_1 and z that can easily be tabulated.

It remains to consider the effect of the second line of Equation 68. This is

$$2i\partial_s\psi_i = \dots - \sum_j \left(\psi_j \int \theta_i^* (\partial_{zz} + 2i\partial_s) \theta_j dS + 2(\partial_z \psi_j) \int \theta_i^* \partial_z \theta_j dS \right).$$

Physically, it is a kind of Coriolis effect that results from a cloud being expanded over changing mode functions. Just as the Coriolis effect can be interpreted as an extra force, these terms can be interpreted as familiar processes from quantum dynamics. The term proportional to ψ_j is equivalent to tunnelling between the components at a rate $\int \theta_i^* (\partial_{zz} + 2i\partial_s) \theta_j dS$. The part $2i \int \theta_i^* \partial_s \theta_j dS$ is atoms being driven from one mode to another by a time-varying repulsion potential; the other part is harder to interpret.

The term in $\partial_z \psi_j$ is less familiar, but it has an obvious interpretation. Atoms in mode j move along the trap with velocity $\partial_z \psi_j$, and therefore the transverse potential they suffer varies with time, at a rate proportional to $(\partial_z \psi_j)(\partial_z \theta_j)$. This drives them into other modes.

When the atoms are confined to the lowest few modes, most of the tunnelling terms can be eliminated adiabatically. According to Equation 68, the phase of the wave function ψ_j rotates at a rate around k_j^2 . Suppose this is much faster than the tunnelling rates, and the energies of the occupied modes from which atoms can tunnel into the empty mode j . Then the tunnelling term will average to zero over each cycle of ψ_j , and no atoms will ever occupy mode j . This is a sufficient condition for the one-dimensional formalism to be valid. It is also a necessary condition, in the sense that if the mode functions change so sharply that atoms scatter into many transverse modes, it would be misleading to describe the dynamics as one-dimensional.

The next step is to add quantum noise to the one-dimensional order parameter. When we do this, the empty modes will have Gaussian random order parameters. However, the argument about separation is still valid, the high-energy modes that initially hold vacuum noise will never be occupied by any atoms, and calculations need only consider the low energy modes. This result, that quantum noise added to energetic transverse modes does not propagate to the occupied modes with low energy, is well known empirically from truncated Wigner simulations.

34. Sampling for Hartree modes

Section 33 explained how the dynamics of a classical order parameter can be computed by expanding it as a one-dimensional field with multiple components for different transverse modes. Computing these classical dynamics is one of the requirements that must be met in order to simulate the dynamics of a quantum gas using the truncated Wigner method. The other requirement is that an initial condition for the classical field can be obtained by sampling the Wigner distribution for the state of the quantum field. This is somewhat harder to achieve, because the transverse modes form part of a mean field theory, not a quantum field theory.

It is possible to proceed as follows. Suppose that the field is initially in a multimode coherent state, with order parameter $\psi(0, x, y, z)$. This mean field has a density $\rho = |\psi|^2$, so Hartree modes $\theta_i(0, x, y, z)$ can be derived as the eigenstates of Equation 66, and the order parameter can be expanded in the form of Equation 64. Supposing that the atomic cloud lies between $-L/2$ and $L/2$ along the z axis, the components ψ_j can be expanded as

Fourier series to give

$$\psi(0, x, y, z) = \frac{1}{\sqrt{L}} \sum_{jn} L^{-\frac{1}{2}} \alpha_{jn} \theta_j(0, x, y, z) e^{2\pi i n z / L}.$$

The functions $L^{-\frac{1}{2}} \theta_j(0, x, y, z) e^{2\pi i n z / L}$ are an orthonormal basis, so the quantum field can be expanded similarly, with

$$\hat{\Psi}(0, x, y, z) = \frac{1}{\sqrt{L}} \sum_{jn} a_{jn} \theta_j(0, x, y, z) e^{2\pi i n z / L}.$$

In the initial state with order parameter $\psi(0, x, y, z)$, the state of the mode a_{jn} is coherent with amplitude α_{jn} . The Wigner distribution of the field can be sampled in the usual way, taking a mean field

$$\psi(0, x, y, z) = \frac{1}{\sqrt{L}} \sum_{jn} (\alpha_{jn} + \xi_{in}) \theta_j(0, x, y, z) e^{2\pi i n z / L},$$

where ξ_{in} is a complex Gaussian variable with $\langle \xi_{in} \rangle = 0$ and $\langle |\xi_{in}|^2 \rangle = \frac{1}{2}$. This has the effect of adding the usual Gaussian noise to the one-dimensional fields ψ_j .

The next step is to integrate the mean field equations of motion given in Equation 69 with the Wigner initial condition, taking spectral derivatives on a finite grid as is usually done in truncated Wigner simulations. For systems where a one-dimensional treatment is physically appropriate, only a few of the components ψ_j will need to be included in the simulation; the other transverse modes will remain in vacuum states.

The final step is taking averages to calculating expectation values. Each Wigner sample generates a set of fields $\psi_j(s, z)$, which are interpreted as if they are samples from the Wigner distribution of the quantum field

$$\hat{\Psi}(s, x, y, z) = \sum_j \hat{\psi}_j(s, z) \theta_j(s, x, y, z).$$

The goal is to find the expectation value of some observable \hat{A} of the field, that can be expanded in the same form as

$$\hat{A}(x, y, z) = \sum_j A_j(s, z) \theta_j(s, x, y, z).$$

The modes θ_j determine the correction from putting this observable into normal order, and in principle these are different for every sample. However, they are unlikely to be very different, and the corrections can be averaged in the same way as the moments. For example, the atomic density could be determined as

$$\rho(s, x, y, z) = \frac{1}{N} \sum_n \left(\left| \sum_j \psi_{nj} \theta_{nj}(s, z) \theta_{nj}(s, x, y, z) \right|^2 - \frac{1}{2} \sum_j |\theta_{nj}|^2 \right),$$

where the outer sum is taken over N samples, each with its own ψ_{nj} and θ_{nj} .

The above derivation assumes that the initial state of the field is coherent. However, the results in Sections 28 and 30 suggest that a coherent state is not an adequate approximation to the ground state of a bosonic gas, at least when density and phase correlations are being considered. If the Hartree modes were going to be used to simulate the splitting of a trap, it would be useful to consider how to initialise them in order to sample the Wigner distribution of a Bogoliubov state instead of a coherent state.

35. Sound wave dynamics

The first theoretical analysis of the spit trap was by Takuya Kitagawa [83], using a Luttinger liquid model. This was invented to describe one-dimensional clouds of electrons, interacting through charge and spin [84]. It is usually solved using Green's functions, and it can be shown to be a universal behaviour of long-wavelength motion in one-dimensional quantum systems [85]. This section will take a different approach that is easier to compare to phase space methods. The dynamics of sound waves in an order parameter are an extension of sound waves in a wave function. We will start with the wave function, following Dirac's approach [26], then consider the effects of atomic repulsion.

In dispersion units, Schrödinger's equation is

$$2i \frac{\partial \psi}{\partial s} = (-\nabla^2 + K^2)\psi.$$

For this section, Leibniz notation will turn up to be more convenient. The wave function can be written in polar form,

$$\psi = A e^{i\phi}, \quad (71)$$

where A and ϕ are real. Now suppose the wave function is a wave packet, where $e^{i\phi}$ is a carrier wave, and A is an envelope that varies gradually over many wavelengths of the carrier. In this limit, the phase ϕ satisfies a Hamilton-Jacobi equation

$$-\frac{\partial \phi}{\partial s} = (\nabla \phi)^2 + K^2,$$

so the dynamics become classical. The probability density A^2 is determined by ϕ , and satisfies an equation of continuity

$$\frac{\partial^2 A}{\partial^2 s} = -2 \frac{\partial}{\partial z} \left(A^2 \frac{\partial \phi}{\partial z} \right). \quad (72)$$

This implies that a particle at position z has velocity

$$v(z) = \frac{\partial \phi}{\partial z}.$$

In geometric units, velocity is identified with the de Broglie wavenumber of a moving atom, and has dimensions of inverse length. For a free particle, the solution to the Hamilton-Jacobi equation is a wave travelling with the velocity of the particle: in the limit of broad

wave packets, the wave packet propagates as a unit, without expanding. This is an entirely classical limit, where the particles follow classical laws of motion, and their density is determined by continuity.

A sound wave model for the order parameter of a Bose-Einstein condensate can be derived similarly [1], starting from Equation 69, the one-dimensional Gross-Pitaevskii equation. The order parameter is written as $\psi = Ae^{i\phi}$, where A and ϕ are real; the density of atoms is $|\psi|^2 = A^2$. The wave packet assumption means that A varies slowly and gradually, so that derivatives of A are small compared to those of ϕ . In a dense condensate, the atoms repel each other strongly density fluctuations are suppressed, so this is a good assumption. Furthermore, a uniform atomic field implies that the atomic velocity $\partial_z\phi$ is uniform, so that higher order derivatives of ϕ are comparable to derivatives of A .

Substituting the form $\psi = Ae^{i\phi}$ in Equation 18 gives

$$i\frac{\partial A}{\partial s} - A\frac{\partial\phi}{\partial s} = e^{-i\phi} \left(-\frac{\partial^2}{\partial z^2} + K_z^2 + gA^2 \right) e^{i\phi} A.$$

Also,

$$e^{-i\phi} \frac{\partial}{\partial z} e^{i\phi} = \frac{\partial}{\partial z} + i\frac{\partial\phi}{\partial z},$$

so the right-hand side reduces to

$$\left(-\left(\frac{\partial}{\partial z} + i\frac{\partial\phi}{\partial z} \right)^2 + K_z^2 + gA^2 \right) A.$$

Taking the real part and neglecting a small term gives

$$\frac{\partial\phi}{\partial s} = -\left(\left(\frac{\partial\phi}{\partial z} \right)^2 + K_z^2 + gA^2 \right). \quad (73)$$

Except for the dependence on A , this is a Hamilton-Jacobi equation. The particles travel classically, and their repulsion is averaged over their density.

Taking imaginary parts gives

$$\frac{\partial A}{\partial s} = -\left(\frac{\partial^2\phi}{\partial z^2} \right) A + 2\frac{\partial\phi}{\partial z} \frac{\partial A}{\partial z},$$

which reduces to Equation 72, so that order parameters satisfy the same continuity equation as do wave functions. Substituting into Equation 73 gives

$$\frac{\partial^2\phi}{\partial s^2} = -\frac{\partial}{\partial s} \left(\frac{\partial\phi}{\partial z} \right)^2 + g \left(2A \frac{\partial A}{\partial z} \frac{\partial\phi}{\partial z} + A^2 \frac{\partial^2\phi}{\partial z^2} \right).$$

The first term is small by assumption; it is unclear why the second term is small. However, neglecting these leaves

$$\frac{\partial^2\phi}{\partial s^2} = gA^2 \frac{\partial^2\phi}{\partial z^2}, \quad (74)$$

which is simply a wave equation. We can assume that A is constant in this equation, with the density of the trap ground state. In this case, the solutions are waves propagating at speed $c_s = 4\sqrt{\pi a_1 A^2}$. In Equation 73, A oscillates rapidly, and can not be assumed constant. The solutions for ϕ determine the velocity of the atoms. If their density is known at an initial time, its values at later times are determined by continuity.

The original derivation of the sound wave model is restricted to one dimension, so it could not simulate the bisection of the condensate. However, the Hamilton-Jacobi equation is not limited to one degree of freedom, and it would be possible to model quantum noise during bisection with three-dimensional sound waves, substituting the mean field in Equation 74 to give a variable speed of sound. However, this might not be much easier than a truncated Wigner simulation with more realistic dynamics given by the Gross-Pitaevskii equation.

Suppose the solution to Equation 74 has the form of small sound waves in a uniform density N_0 . Then Equation 72 can be approximated as

$$\frac{\partial A^2}{\partial s} = -2N_0^2 \frac{\partial^2 \phi}{\partial z^2}.$$

This is now linear in A^2 and ϕ , so sound waves obey a superposition principle, provided their amplitudes are small. This has a solution

$$\phi = \epsilon_k \cos(kz) e^{i\chi s}, \quad A^2 = N_k \cos(kz) e^{i\chi s},$$

provided

$$\frac{N_k}{N_0} = -\frac{2ik}{c_s} \epsilon_k.$$

Note that the real part of ϵ_k determines the phase ϕ , and its imaginary part determines the density N_k . The small amplitude assumption is that $N_1/N_0 \approx \epsilon \ll 1$. There exist similar solutions with sines in place of the cosines, and the general solution is

$$\phi = \sum_{k>0} (\epsilon_k \cos(kz) + \epsilon'_k \sin(kz)) e^{i\chi s}, \quad \frac{A^2}{N_0} = 1 - \sum_{k>0} \frac{2ik}{c_s} (\epsilon_k \cos(kz) + \epsilon'_k \sin(kz)) e^{i\chi s}.$$

Every expression below of the form $A \cos kz$ implies an additional term $A' \sin kz$.

The wave equation is second order, so its initial condition must specify ϕ and $\frac{\partial \phi}{\partial s}$. The goal is to choose these to be equivalent to sampling ψ from a Wigner distribution. To determine them, the sound wave solution must be substituted back into $\psi = A e^{i\phi}$. The solution A^2 to Equation 74 is complex, which is allowed because the equation is linear. To find A , however, the real part of A^2 must be taken explicitly. For the initial conditions, $s = 0$, so

$$\frac{A^2}{N_0} = 1 + \sum \frac{2k}{c_s} \text{Im} \epsilon_k \cos(kz),$$

and to first order in ϵ_k ,

$$\frac{A}{\sqrt{N_0}} = 1 + \sum \frac{k}{c_s} \text{Im} \epsilon_k \cos(kz).$$

Substituting this, and keeping the first order terms, gives

$$\begin{aligned}\psi &= \sqrt{N_0} \left(1 + \sum \frac{k}{c_s} \text{Im } \epsilon_k \cos(kz) \right) \left(1 + i \sum \text{Re } \epsilon_k \cos(kz) \right) \\ &= \sqrt{N_0} + \sum \sqrt{N_0} \left(\frac{k}{c_s} \text{Im } \epsilon_k + i \text{Re } \epsilon_k \right) \cos kz.\end{aligned}$$

For normalisation purposes, suppose that the condensate extends over a length L , and wherever $\cos kz$ appears in the last two paragraphs it means $\sqrt{2/L} \cos(kz)$. When ψ is drawn from the Wigner distribution for a coherent state, the variances of the real and imaginary coefficients of $\cos kz$ are both $\frac{1}{4}$. This can be achieved by drawing ϵ_k with variances

$$V(\text{Re } \epsilon_k) = \frac{1}{4N_0} = \frac{k^2}{c_s^2} V(\text{Im } \epsilon_k).$$

This corresponds to the variance of ϕ being $1/N_0$, and of N being N_0 , as expected for a coherent state. The quantity k is actually an atomic wave number; c_s is the wave number of an atom moving at the speed of sound. Therefore, k/c_s is a kind of Mach number for the atoms.

For a mode of the atomic field to be a coherent state, the corresponding mode of the Luttinger liquid must be in a squeezed state. For subsonic modes, with $k < c_s$, the quadrature corresponding to the phase of the field is squeezed, and that corresponding the atom number is anti-squeezed. As time passes, the squeezed quadrature rotates, and the phase of the field is scrambled.

These sound wave dynamics approximate the mean-field dynamics. If noise is added as above, the result will approximate the samples of a truncated Wigner simulation. When the condensate is bisected, it can be treated as passing through a beam splitter, on the assumption that the bisection is instant. In this case, the twin parts can be given the same mean field, with independent quantum noise added. The mean field will contain thermal fluctuations from the original condensate, but this is common to both bisected parts.

Because this is a linear wave equation, the solution obeys a superposition principle. It is therefore clear that the common and relative phases of two uncoupled condensates evolve independently at this level of approximation. This will presumably not be the case with mean field dynamics, where fluctuations in the common density propagate into the relative phase through the nonlinear term $\dot{\psi} \propto N\psi$.

36. Field theory

Sections 32 and 33 discussed how to expand an order parameter over transverse modes that adapted to the repulsive field. This expansion could be used to simulate the dynamics in a truncated Wigner model. However, this could be done more rigorously if the quantum field was expanded instead of the order parameter, to reduce a three-dimensional quantum field $\hat{\psi}(x, y, z)$ to a system of one-dimensional fields $\hat{\Psi}_j(z)$.

One way to proceed is to make the transverse modes observables, so that $\hat{\theta}_j(x, y, z)$ is a function of the $\hat{\psi}_j(z)$, but not of the field operators $\hat{\psi}_j(z')$ at other points on the axis. More

precisely, $\hat{\theta}_j(z)$ should be a function of the atomic density $\hat{n}(z)$, whose value is the solution to Equation 64 for a given density.

In second quantisation, the Hamiltonian for an atomic field is given by Equation 14,

$$\hat{H} = \int \hat{\Psi}^\dagger(r)(-\nabla^2 + K^2(r))\hat{\Psi}(r) d^3r + 8\pi a \int \hat{\Psi}^\dagger(r)\hat{\Psi}^\dagger(r)\hat{\Psi}(r)\hat{\Psi}(r) d^3r.$$

In a quantum field theory, the atomic density is an observable $\hat{n}(r) = \hat{\Psi}^\dagger(r)\hat{\Psi}(r)$, and the transverse modes $\theta_i(r, \hat{n}(r))$ are functions of this. This is to say, the modes are complex fields when the atomic field is in a density eigenstate, but in general different atomic densities are superposed. The modes can still be defined as solutions to the eigenvalue problem

$$(-\partial_{xx} - \partial_{yy} + K^2 + 8\pi a\hat{n})\theta_i = (\hat{k}_i)^2\theta_i, \quad (75)$$

but these eigenvalues are operators themselves, that depend on the operator \hat{n} .

Given a fixed eigenstate $|n(r)\rangle$, the modes are complete, and satisfy the usual orthonormality condition,

$$\int \theta_i(n, x, y, z)\theta_j(n, x, y, z) dS = \delta_{ij}, \quad (76)$$

so the field operator $\hat{\Psi}(x, y, z)$ can be expanded as follows:

$$\begin{aligned} \hat{\Psi}(x, y, z) &= \sum_{nn'} |n\rangle\langle n| \hat{\Psi}(x, y, z) |n'\rangle\langle n'| \\ &= \sum_{inn'} |n\rangle \left(\theta_i(n, x, y, z) \int \theta_i^*(n, x', y', z) \langle n| \hat{\Psi}(x', y', z) |n'\rangle dx' dy' \right) \langle n'| \\ &= \sum_i \theta_i(\hat{n}, x, y, z) \left(\int \theta_i^*(\hat{n}, x', y', z) \hat{\Psi}(x', y', z) dx' dy' \right) \\ &= \sum_i \theta_i(\hat{n}, x, y, z) \hat{\Psi}_i(z). \end{aligned} \quad (77)$$

There is also an orthonormality condition for the operators $\theta_i(\hat{n}, x, y, z)$,

$$\begin{aligned} \int \theta_i(\hat{n}, x, y, z)\theta_j(\hat{n}, x, y, z) dx dy &= \sum_n \int \theta_i(n, x, y, z) |n\rangle\langle n| \theta_j(n, x, y, z) dx dy \\ &= \sum_n \delta_{ij} |n\rangle\langle n| = \delta_{ij}. \end{aligned}$$

If the expansion of Equation 77 is substituted into the single-particle term of the Hamiltonian, and Equation 75 applied, the result is approximately

$$\int \hat{\Psi}^\dagger(-\nabla^2 + K^2)\hat{\Psi} = \sum_i \int \hat{\Psi}_i^\dagger(z)(-\partial_{zz} + \hat{k}_i^2)\hat{\Psi}_i(z) dz - 8\pi a \int \hat{\Psi}^\dagger \hat{n} \hat{\Psi}.$$

One approximation involved is that $\theta_i(\hat{n})$ commutes with $\hat{\Psi}$, which is true to the extent that $\theta_i(n+1)$ is close to $\theta_i(n)$ —physically, the condition is that adding one atom has little effect

on the mode functions. The other approximation is that θ_i varies gradually with z . The second term exactly cancels the two-particle term of the Hamiltonian, leaving

$$\hat{H} = \sum_i \int \hat{\Psi}_i^\dagger(z) (-\partial_{zz} + \hat{k}_i^2) \hat{\Psi}_i(z) dz.$$

Writing

$$\hat{k}_{0i}^2 = \int \theta_i^* (-\partial_{xx} - \partial_{yy} + K^2) \theta_i dx dy,$$

and commuting θ_i and \hat{k}_i^2 , both being functions of \hat{n} , this becomes

$$\hat{H} = \sum_i \int \hat{\Psi}_i^\dagger(z) \left(-\partial_{zz} + \hat{k}_{0i}^2 + 8\pi a \int \theta_i^* \hat{n} \theta_i dS \right) \hat{\Psi}_i(z) dz. \quad (78)$$

Chapter 5 Expansions over Coherent States

Quantum and atom optical systems include many particles distributed among many modes. Their quantum mechanical states are vectors in Hilbert spaces of high dimension, far too high for the components to be stored on any possible computer. In Chapters 3 and 4, these states were represented by randomly sampling amplitudes in phase space. When the fair sampling assumption holds, this method makes the impossible quite easy. When it doesn't, the impossible fails to converge.

Although the simplest mechanical systems allow Schrödinger's law of motion to be solved in closed form, most problems in quantum mechanics must be solved approximately. This is often done using variational methods. These specify a class of approximate solutions that can be manipulated algebraically, then identify the member of that class which most nearly satisfies the laws of motion. Coherent states, or Gaussian wave packets, have often been used as variational approximants, because they are defined by their width and their central position and momentum, making them a simple and general approximant to any localised quantum state.

Much work has gone into a class of approximants that are called, paradoxically, multi-configurational Hartree states [86]. These approximants have been applied to atom optics, but they tend to require very many parameters to represent the many-particle states involved.

This chapter describes another approach. Instead of sampling coherent amplitudes at random, we will hope to expand the state over a small set of coherent states, and have the coherent and probability amplitudes evolve in time so that their superposition approximates the state of the system. This is a type of variational method. For reasons that this chapter elucidates, this procedure is only practical when the state overlaps with a small region of quantum phase space.

This method might sound very similar to the phase space methods discussed in Chapter 3. In the P and positive P methods, the density operator of the system is expanded over coherent state projectors as $\rho = \int P(\alpha^*, \beta) |\alpha\rangle\langle\beta| d^2\alpha d^2\beta$. The method presented in this Chapter applies to kets instead of density operators, but it could be modified to apply to a density operator. In that case, the variational ansatz would be $\rho = \sum_i P_i |\alpha_i\rangle\langle\beta_i|$, a very similar form to the phase space representations.

Although the coherent state expansion is very similar to the existing phase space representations, the variational dynamics are quite different from the Monte Carlo dynamics

that have traditionally been used with these representations. In the Monte Carlo methods, each trajectory (α_i, β_i) evolves independently of the other trajectories. In the variational methods, the whole ensemble of coherent amplitudes is evolved as one set. In return for this more complex implementation, the dynamics should converge to quantum dynamics in the limit of a large ensemble, and there is no stochastic term in the equations of motion, avoiding the large sampling errors that prevent positive P simulations from continuing past a certain time.

This way of discretising phase space was originally suggested by Heller. In the 1970s, he put variational mechanics and phase space expansions together [14], and proposed to simulate quantum dynamics variationally, using superpositions of a finite set of coherent states as the approximants. At first blush, this approach seems very sensible. In the limit of a large set of coherent states, the finite superposition converges to a continuous phase space expansion, and every quantum state can be approximated in this way. The experience of quantum optics shows that there are interesting systems in which this converges while the number of states in the set is reasonably small, and variational dynamics should avoid the convergence problems that limit the Monte-Carlo methods.

There have been several attempts to implement Heller's proposal [57, 18], but none of them has entirely succeeded. In all cases, numerical instability has prevented the superposition being expanded to the point of convergence. These attempts have aimed to solve chemistry problems that involve dozens of atoms. Their failure to converge limited their precision, but useful results were obtained from superpositions with a few components, and ad-hoc methods to improve their stability. For example, the trajectories of the centres of the wave packets can be constrained to be classical trajectories [14]; this makes the calculations more stable; this will share some of the limitations of the truncated Wigner approach, for example it must fail when the wave function expanded by the packets tunnels through a barrier or does anything else inconsistent with classical mechanics.

The problems that are tackled with phase space methods come from atom optics or condensed matter; even in second quantisation, they involve thousands of modes. The instability would prevent any useful information being obtained about these problems, so a systematic approach to regularisation is required in order to tackle these problems variationally. The prospective rewards are significant: one problem for which Monte Carlo methods struggle involves metastable quantum fields that are relevant to the dynamics of the early universe [87].

The physicists and quantum chemists who encountered the instability of variational Gaussian methods did not systematically investigate their origin. However, work in other fields sheds some light: expanding a function as a sum of Gaussian wave packets is a somewhat obvious idea, which has occurred to mathematicians and signal engineers as well as to physicists. The research effort that developed wavelet expansions studied Gaussian wave packets too [2]. What physicists call a phase space expansion, these mathematicians call a windowed Fourier transform. They have proved some general results about discrete sets of coherent states. They have also studied overcomplete sets of vectors, and the stability of

expansions over them, with results that were introduced in Section 12. Many of the general ideas offer insight into the specific case of coherent state frames.

There is another source of instability that stems from the dynamics of Schrödinger's equation. This occurs whenever Schrödinger's equation is integrated with a quartic Hamiltonian, over discrete time steps. In the second quantised picture, the repulsion between particles causes the amplitudes to oscillate very rapidly where many particles occupy a mode. The simulation becomes stable when these components are removed, but that changes the dynamics. Integrating Schrödinger's equation with repulsive particles either requires very short time steps, or special integration formulae that exploit features of unitary dynamics [88]. It is difficult to apply these formulae to the variational dynamics of the coherent states, because their dynamics are not unitary.

This chapter will start by discussing expansions of kets as finite sums of coherent states, then it will discuss the implications for dynamical simulations. Some insight is gained into the stability problems. They are not fully resolved, but some steps for further development are proposed.

37. Multivariate derivatives

This thesis will use a modern notation for functions and partial derivatives, based on the one popularised by Spivak [89]. We will extend it to treat functions whose values are kets and operators, not only real and complex numbers. The thesis will consider complicated least squares problems, and this notation will allow them to be described concisely.

Suppose that a certain atomic potential has a ground state $\phi(x, y, z)$, and an atom that moves through the potential has the wave function $\psi(t, x, y, z)$. We will write the partial derivatives of these functions in a compact notation, for example

$$D_2\phi(x_0, y_0, z_0) = \left. \frac{\partial\phi}{\partial y} \right|_{x=x_0, y=y_0, z=z_0}.$$

The partial derivative $D_2\phi$ is a function similar to ϕ . The subscript 2 indicates a partial derivative with respect to the second argument. For functions such as ϕ , whose arguments are spatial coordinates, the count starts from 1. For functions of time such as ψ , the time will be the first argument, and the notation ∂_t will be used for the partial derivative. So in this case

$$D_2\psi = \frac{\partial\psi}{\partial y} \quad \text{and} \quad \partial_t\psi = \frac{\partial\psi}{\partial t}.$$

Multiple derivatives are written as

$$D_{32}\phi = \frac{\partial^2\phi}{\partial z\partial y},$$

and Laplacian operators over some of the spatial axes as

$$\nabla_n^2 f = D_{11}f + \cdots + D_{nn}f.$$

The three-dimensional Laplacian ∇_3^2 is also written simply as ∇^2 .

The partial derivatives compose a Jacobian matrix

$$\phi' = (D_1\phi \quad D_2\phi \quad D_3\phi).$$

The Jacobian ϕ' is a function of the coordinates. The value $\phi'(x, y, z)$ is a row vector of three complex numbers, and it satisfies

$$\phi(r + h) = \phi(r) + \phi'(r) \cdot h + o(h^2), \quad (79)$$

where r and h are 3-vectors. The arguments of a function will sometimes be identified with a column vector, as in the matrix product $\phi'(r) \cdot h$. If, instead of the scalar function ϕ , we considered a vector-valued function, perhaps a trapping magnetic field $B(x, y, z) = (B^1(x, y, z), B^2(x, y, z), B^3(x, y, z))$, the Jacobian would still be a function of the coordinates, but the value $B'(x, y, z)$ would be a 3×3 matrix instead of a 1×3 row vector:

$$B' = \begin{pmatrix} D_1B^1 & D_2B^1 & D_3B^1 \\ D_1B^2 & D_2B^2 & D_3B^2 \\ D_1B^3 & D_2B^3 & D_3B^3 \end{pmatrix},$$

However, Equation 79 is still satisfied with B in place of ϕ . In this notation, the chain rule can be written

$$(f \circ g)' = (f' \circ g)g',$$

the same form as in ordinary calculus. The difference is that the derivative functions are matrix-valued, and the right hand side is a matrix product.

In the case of analytic functions of a complex variable, the partial derivatives are the usual complex derivatives. For more general complex functions, we could define derivatives D_i^* as in Wirtinger calculus. The work presented in this thesis does not require such derivatives, but some of the extensions proposed in it do.

A useful convention is that nested vectors concatenate. For example, if $r = (x, y, z)$, then $(t, r) = (t, x, y, z)$.

We will take derivatives of functions whose values are kets and operators, not just real and complex numbers. An example of a ket valued function is $|f(z, w)\rangle = z|A\rangle + w^2|B\rangle$. The partial derivatives $|D_1f(z, w)\rangle = |A\rangle$ and $|D_2f(z, w)\rangle = 2w|B\rangle$ are ket-valued functions similar to $|f\rangle$. The derivative operator D_i is written inside the ket because it acts on the function, not on the ket that is the value of the function.

In this notation, Schrödinger's equation is written

$$i\hbar|\Psi'(t)\rangle = H|\Psi(t)\rangle. \quad (80)$$

Here, $|\Psi\rangle$ is a function of time, whose value at time t is the ket $|\Psi(t)\rangle$. The derivative of the function $|\Psi\rangle$ is the function $|\Psi'\rangle$. In contrast, the Hamiltonian H acts on the value $|\Psi(t)\rangle$, so it is written outside the ket. Schrödinger's equation could be abbreviated to

$$i\hbar|\Psi'\rangle = H|\Psi\rangle,$$

equating the functions instead of their values at time t . Since $|\Psi\rangle$ is a function of a single variable, the Jacobian is a scalar $|\Psi'\rangle = |\partial_t\Psi\rangle$, and Equation 80 can be written

$$i\hbar|\partial_t\Psi\rangle = H|\Psi\rangle.$$

38. Notation for kets and vectors

Quantum mechanics is formally based on a Hilbert space of bras and kets, but our computers must work in finite dimensional spaces whose elements expand to matrices and column vectors. In order to describe numerical quantum mechanics, this chapter combines Dirac notation with matrices, using a sort of block matrix in which some blocks are matrices of complex numbers while others are elements of Hilbert spaces.

A linear combination of kets has the form $|\psi\rangle = c_1|a_1\rangle + c_2|a_2\rangle + \cdots + c_r|a_r\rangle$. This can be written as a matrix product $|\psi\rangle = |A\rangle c$, where the kets form a row vector $|A\rangle = (|a_1\rangle \cdots |a_r\rangle)$, and their coefficients form a column vector $c = (c_1 \cdots c_r)^T$. When the kets are coherent states, the function $|o\rangle$ defined in Section 4 evaluates to the row vector $|o(\alpha_1, \dots, \alpha_r)\rangle = (|\alpha_1\rangle \cdots |\alpha_r\rangle)$. All of the coherent states in this chapter are of a single mode, so there can be no confusion about their assignment to modes and states.

The vector $|A\rangle$ defines a linear operator taking vectors of complex numbers to kets. Its adjoint $\langle A| = (\langle a_1| \cdots \langle a_r|)^T$ takes kets back to column vectors. Under the rules of matrix multiplication—or the equivalent rules of operator composition—the product

$$\langle A|A\rangle = \begin{pmatrix} \langle a_1|a_1\rangle & \cdots & \langle a_1|a_r\rangle \\ \vdots & \ddots & \vdots \\ \langle a_r|a_1\rangle & \cdots & \langle a_r|a_r\rangle \end{pmatrix}$$

is an operator on \mathbb{C}^r . On the other hand, the product $|A\rangle\langle A| = |a_1\rangle\langle a_1| + \cdots + |a_r\rangle\langle a_r|$ operates on Hilbert space.

Instead of regarding $|A\rangle$ as a $1 \times r$ vector, we could regard the kets as columns of infinite length, making $|A\rangle$ an $\infty \times r$ matrix and the operator $|A\rangle\langle A|$ an $\infty \times \infty$ matrix.

Other matrix products are possible. Provided the kets $|a_i\rangle$ are linearly independent, there will be a singular value decomposition as discussed in Section 10. This will have the form $|A\rangle = |U\rangle S V^\dagger$, where $|U\rangle$ is a row vector of r orthonormal kets, S is a diagonal $r \times r$ matrix of complex numbers, and V is a unitary $r \times r$ matrix. In fact, $|A\rangle$ has an SVD when the kets $|a_i\rangle$ are linearly dependent, but $|U\rangle$ and S are smaller in this case, and the span of $|U\rangle$ has dimension less than r . In order to make up r orthonormal kets, some extra dimensions would need to be chosen from Hilbert space, and there is no consistent way to select them. Instead of using the full form of the SVD, we will use the reduced form where $|U\rangle$ is a row of m kets and S is a $m \times r$ matrix; here m is the rank of $|A\rangle$.

As well as mappings from kets to vectors, we will also consider linear mappings from kets to infinite dimensional spaces. Computer simulations of fields store discrete approximations to them, and the continuous field is often simpler to analyse than the discretisation. Let $\langle\Phi|$ be a linear mapping from kets to functions of a complex variable. This defines a

linear pointwise mapping $\langle \Phi(\alpha) |$ from kets to complex numbers, where $\langle \Phi(\alpha) | \psi \rangle = f(\alpha)$ and $f = \langle \Phi | \psi \rangle$. Therefore mappings from kets to functions of a complex variable can be identified with functions taking complex numbers to bras. We have already seen one of them: the function $\langle o |$ that takes complex amplitudes to coherent states also takes kets to phase space functions, where we consider an inner product of a function and a ket such as $\langle o | \psi \rangle$ to be a function itself.

Going the other way, we can identify a ket-valued function such as $|o\rangle$ with a conjugate-linear mapping that takes the function f of a complex variable to the corresponding expanded ket $|\psi\rangle = \int f^*(\alpha) |o(\alpha)\rangle d^2\alpha$. This integral has the same form as the L^2 inner product on functions, where the inner product of f and g is $\int f^*g$.

Considered in this way as linear mappings, functions such as $|\Phi\rangle$ and $|o\rangle$ have singular value decompositions. These have the form

$$\int f^* |\Phi\rangle = \sum_{i=1}^{\infty} |u_i\rangle \sigma_i \int f^* v_i,$$

where the left singular kets $|u_i\rangle$ are an orthonormal basis for ket space, and the right singular functions v_i are orthonormal under the L^2 inner product. The coherent states form a tight frame for ket space: by Equation 8 every singular value of $|o\rangle$ is $1/\pi$. Therefore its singular kets can be any orthonormal basis; it is often convenient to choose the number states, in which case the right singular functions are $f_n(\alpha) = \pi^{-\frac{1}{2}} \langle \alpha | n \rangle$.

More general matrices are sometimes useful. For instance, the operator

$$M = \begin{pmatrix} 0 & |A\rangle \\ \langle A| & 0 \end{pmatrix}$$

acts on columns of the form $(|\psi\rangle \quad z_1 \quad \cdots \quad z_r)^T$ to give columns of the same form. Indices are ordinal numbers, so $\infty + r$ is a valid size. The top left zero in the above matrix is an operator on Hilbert space, while the bottom right one is an $r \times r$ matrix: the elements in each row must have the same column dimension and visa-versa.

39. How the variational method is supposed to work

This section presents a simple variational method where the approximant is a linear combination of coherent states, in order to demonstrate what we intend to do. This method is too simple to be numerically stable. The following sections elaborate why it is unstable, consider what might be done to improve it, and discuss which kinds of simulated systems are likely to generate stable variational dynamics as a result.

The variational approximants to the state $|\Psi(t)\rangle$ that we consider in this chapter have the form

$$|\psi(\phi, \alpha)\rangle = \sum_{r=1}^R e^{\phi^r + \alpha^r a^\dagger} |0\rangle. \quad (81)$$

Using the notation of the last section, this can be written in the form

$$|\psi(\phi, \alpha)\rangle = |o(\alpha)\rangle c \quad (82)$$

where $|o(\alpha)\rangle = (|\alpha^1\rangle \dots |\alpha^R\rangle)$ and $c = (e^{\phi^1 + \frac{1}{2}|\alpha^1|^2} \dots e^{\phi^R + \frac{1}{2}|\alpha^R|^2})^T$. This form is simpler, but Equation 81 is an analytic function of ϕ and α , which is convenient for optimisation.

In the case that ϕ and α are R -vectors, and $z = (\phi, \alpha)$ is the $2R$ -vector formed by concatenating them, and $|\psi(z)\rangle = |o(\alpha)\rangle c$, we will refer to z as the variational parameters, α as the variational amplitudes, and c as the variational coefficients. In the limit of large R , the sum in Equation 81 becomes an integral over phase space, and this approximant can represent any state according to Equation 9.

The variational approximation supposes that the solution $|\Psi(t)\rangle$ has the form $|\psi(\zeta(t))\rangle$, where $\zeta(t)$ is a vector of time-varying variational parameters. This form can be substituted in Schrödinger's equation,

$$|\Psi'(t)\rangle = -iH|\Psi(t)\rangle. \quad (83),$$

to give

$$|\psi'(\zeta(t))\rangle \cdot \zeta'(t) = -iH|\psi(\zeta(t))\rangle. \quad (84)$$

For the discussion of variational methods, we will set $\hbar = 1$. The Jacobian vector $|\psi'(\zeta(t))\rangle$ on the left hand side is composed of partial derivatives $|D_i\psi(\zeta(t))\rangle$. These are multiplied by coefficients from the column vector $\zeta'(t) = ((\zeta^1)'(t) \dots (\zeta^{2R})'(t))^T$. When a single coherent state is parameterised in the form of Equation 81, its Jacobian is

$$|\psi'(\phi, \alpha)\rangle = (|D_1\psi(\phi, \alpha)\rangle \quad |D_2\psi(\phi, \alpha)\rangle) = (1 \quad a^\dagger) |\psi(\phi, \alpha)\rangle.$$

When the approximant is a superposition of coherent states with $R > 1$, the Jacobian is

$$|\psi'(\phi, \alpha)\rangle = (|\psi(\phi^1, \alpha^1)\rangle \quad \dots \quad |\psi(\phi^R, \alpha^R)\rangle \quad a^\dagger|\psi(\phi^1, \alpha^1)\rangle \quad \dots \quad a^\dagger|\psi(\phi^R, \alpha^R)\rangle). \quad (85)$$

Unless the solution to Schrödinger's equation has the exact form of the approximant, no function ζ will satisfy Equation 84 exactly. The elements of $|\psi'(\zeta(t))\rangle$ span a $2R$ -dimensional subspace of Hilbert space, and the right hand side $H|\psi(\zeta(t))\rangle$ will usually fall outside this subspace. We can, however, find a least-squares solution: if we allow $\zeta(t)$ to be an arbitrary $2R$ -vector z , there is some $2R$ -vector $\dot{\zeta}_H(z)$ that minimises the Hilbert-space residual

$$\| |\psi'(z)\rangle \cdot \dot{\zeta}_H(z) + iH|\psi(z)\rangle \|^2. \quad (86)$$

This is a linear least squares problem, whose solution defines the function $\dot{\zeta}_H$. The system of ordinary differential equations

$$\zeta'(t) = \dot{\zeta}_H(\zeta(t)) \quad (87)$$

can be solved for some variational parameters $\zeta(t)$, and hopefully $|\psi(\zeta(t))\rangle$ approximates the state $|\Psi(t)\rangle$.

The solution to the least squares problem $\dot{\zeta}_H$ can be calculated using the pseudoinverse that was introduced in Section 12. The Jacobian vector $|\psi'(z)\rangle$ has a singular value decomposition

$$|\psi'(z)\rangle = |U\rangle S V^\dagger,$$

and the kets $|u^i\rangle$ that compose $|U\rangle$ are an orthonormal basis for the space of kets that can be expressed as $|\psi'(z)\rangle \cdot x$ for some coefficient vector x . The solution $\dot{\zeta}_H$ is the vector for which $|\psi'(z)\rangle \cdot \dot{\zeta}_H$ is closest to $-iH|\psi(z)\rangle$. These kets are closest when their components along the basis $|u^i\rangle$ are equal, with

$$-i\langle U|H|\psi(z)\rangle = \langle U|\psi'(z)\rangle \cdot \dot{\zeta}_H = SV^\dagger \dot{\zeta}_H.$$

The solution is

$$\dot{\zeta}_H(z) = -i\langle (\psi')^+(z)|H|\psi(z)\rangle, \quad (88)$$

where the pseudoinverse is $\langle (\psi')^+(z)| = VS^{-1}\langle U|$.

The kets $|D_i\psi(z)\rangle$ might not be linearly independent. In this case, as discussed in Section 38, the vector $|U\rangle$ and the matrix V have fewer than $2R$ columns, and there are only that many singular values in the matrix S . The columns of $|U\rangle$ still form a basis for states of the form $|\psi'(z)\rangle \cdot \dot{\zeta}$, the argument remains valid, and Equation 88 still solves the least squares problem.

This process is called the Dirac-Frenkel variational method [90, 91]. It can be used with entirely general forms of $|\psi\rangle$, but this chapter will stick with the form of Equation 81.

All this can be interpreted geometrically. The right hand sketch in Figure 29 depicts a curved manifold in Hilbert space, comprising the states that have the form $|\psi(x)\rangle$ where x lies in the volume of phase space depicted to the left. Suppose that the system currently has state $|\psi(z)\rangle$. The rate of change required by Schrödinger's equation, $-iH|\psi(z)\rangle$, will take the state off the manifold, so that it can not be represented by an approximant of the right form. This is shown as a vector starting from $|\psi(z)\rangle$.

The shaded surface is the tangent space to the manifold, the kets that are $|\psi'(z)\rangle \cdot \dot{\zeta}$ for some ζ . This is the subspace spanned by the kets $|u^i\rangle$. The least squares solution $|\psi'(z)\rangle \cdot \dot{\zeta}$ is the projection of the rate of change $-iH|\psi(z)\rangle$ onto this subspace; the dashed line is normal to the tangent plane. In Figure 29, the correct rate of change is almost normal to the manifold, and so the variational approximation would be very poor. Hopefully, in actual simulations, the rate of change will be almost tangent to the manifold, and the two vectors nearly equal.

The Dirac-Frenkel solution is only optimised locally. There might be some other parameter function w such that, at large times, $|\psi \circ w\rangle$ approximates $|\Psi\rangle$ better than the Dirac-Frenkel solution $|\psi \circ \zeta\rangle$ does. This would require that, at early times, $|D\psi \circ \zeta\rangle$ is closer to $-iH|\psi \circ \zeta\rangle$ than $|D\psi \circ w\rangle$ is to $-iH|\psi \circ w\rangle$, but that ζ loses the ability to track the dynamics at later times. When we discuss regularisation below, we will see that this is in fact the usual case.

In this chapter, we will suppose that the variational method is being used to solve Schrödinger's equation. However, it could be applied to the Liouville equation for a density operator, or to a quantum master equation, with very little modification. Instead of expanding a state $|\psi\rangle$ over Gaussian wave packets $|\alpha_r\rangle$ with dimension equal to the number of modes, we would expand a state ρ over Gaussian wave packets $|\alpha_r\rangle\langle\beta_r|$ with dimension




Figure 29. The Dirac-Frenkel method in parameter space and Hilbert space.

twice the number of modes, and project the Liouvillian superoperator $\mathcal{L}\rho$ onto their tangent space. There is one complication. Every ket $|\alpha\rangle\langle c|$ actually represents a quantum state, but linear combinations of $|\alpha_r\rangle\langle\beta_r|$ are not necessarily density operators. Combinations of the form $|\alpha_r\rangle\langle\alpha_r|$ are all density operators, but this is a discretisation of the P distribution, so it has the limitations discussed in Section 16 and in [40].

Having described how the method is supposed to work, we now turn to how it fails. Equation 88, the Dirac-Frenkel equation of motion, must usually be solved with a finite difference formula whose stability depends on the Jacobian of $\dot{\zeta}_H(z) = \langle(\psi')^+(z)|H|\psi(z)\rangle$. The mapping $|\psi'\rangle$ can be badly conditioned, in the case where the partial derivatives $|D_i\psi(z)\rangle$ do not form a tight frame for some subspace of Hilbert space that includes $H|\psi(z)\rangle$. In this case, the Jacobian can have large eigenvalues, so that Equation 88 will require an impractically small time step to solve numerically.

The next sections will explain how, for many states, the expansion of Equation 81 requires very large variational coefficients to give the best approximation, and this causes the linear transformation ζ' to become badly conditioned. The following sections will explain how this causes Equation 87 to become stiff, in the sense introduced in Section 11, and require impractically small time steps to solve. Finally, starting in Section 44, we will discuss some ways to improve the conditioning while maintaining variational accuracy.

40. Expansions over frames

Section 39 derived Equation 87, the variational equation of motion, and concluded that it would be stiff unless the kets $|\psi'(z)\rangle$ form a snug frame for some ket subspace that includes $H|\psi(z)\rangle$. Section 13 analysed grids of sinc wave packets as a frame for functions; the analysis of coherent states as a frame for wave functions is quite similar, because the wave function of a coherent state is a Gaussian wave packet, but we will carry it out using

Dirac notation.

The Jacobian $|\psi'(\phi, \alpha)\rangle$ in Equation 85 includes two types of element. The partial derivatives $|D_{R+1}\psi\rangle$ through $|D_{2R}\psi\rangle$ have the form $a^\dagger|\psi(\phi^r, \alpha^r)\rangle$. We will defer considering these until Section 43. For the next sections, we will consider expansions over the kets $|D_1\psi\rangle$ through $|D_R\psi\rangle$. These are unnormalised coherent states $|\psi(\phi^r, \alpha^r)\rangle$, where the normalisation is the variational coefficient $c(\phi^r, \alpha^r) = e^{\phi^r + \frac{1}{2}|\alpha^r|^2}$. The varying sizes of the elements affect the condition of the frame as discussed in Section 12. When $H|\psi(z)\rangle$ is expanded, the small kets will be assigned large coefficients, which will cause the equation of motion to be stiff in a fairly obvious way. For actual computation, it would be best to use the normalised parameters of Equation 82, expand everything over real and imaginary parts, and avoid this. But this leaves the question of the condition of the normalised frame $|o(\alpha)\rangle$. This is quite interesting, and is treated in the next sections.

There are two reasons to be concerned about the condition of expansions over a discrete set of coherent states. The first issue is that a finite set of variational amplitudes α_i is contained in a finite volume of phase space. The state $|\psi\rangle$ being expanded might have components $\langle\beta|\psi\rangle$ where the amplitude β is distant from this volume, and thus by Equation 7 the state $|\beta\rangle$ is nearly orthogonal to every frame element $|\alpha_i\rangle$. The coefficients that extrapolate this component of $|\psi\rangle$ have to be large and oscillating, so that the components of different $|\alpha_i\rangle$ nearly cancel out. Section 41 considers this issue.

The second issue is how densely the sample amplitudes fill the volume that contains them. The situation here is very similar to the sinc interpolants that were analysed in Section 13: there is a critical sampling density, below which a discrete grid of coherent states fails to span ket space, and the condition of the frame gets rapidly worse as the density decreases towards the limit. This is discussed in Section 42.

The formalism for expansions over continuous volumes of quantum phase space differs somewhat from that for discrete expansions. A discrete expansion has a vector of coefficients c , which the expansion operator $|o(\alpha)\rangle$ maps to the expanded state $|o(\alpha)\rangle c$. When the full quantum phase space is used, the coefficients are a function of a complex variable $f(\alpha)$, which the expansion operator $|\Gamma\rangle$ maps to a ket as

$$|\Gamma\rangle f = \int f(\alpha)|\alpha\rangle d^2\alpha.$$

It follows from Equation 8 that

$$\|\langle\Gamma|\psi\rangle\|^2 = \int |\langle\alpha|\psi\rangle|^2 d\alpha = \langle\psi|\int |\alpha\rangle\langle\alpha| d^2\alpha|\psi\rangle = \pi,$$

so the full quantum phase space is a tight frame for Hilbert space, with bound π .

This holds more generally than for coherent states. It is the case for any normalisable wave packets $|\psi_0\rangle$ that

$$\int d^2\alpha D^\dagger(\alpha)|\psi_0\rangle\langle\psi_0|D(\alpha) = \pi, \quad (89)$$

where $D(\alpha)$ is the phase space displacement operator [2].

The operator $|\Gamma\rangle$ has a singular value decomposition, as discussed in Section 38. Therefore it has a pseudoinverse $\langle\Gamma^+|$, and expansion functions can be constructed as $f = \langle\Gamma^+|\psi\rangle$. This defines a linear functional $\langle\Gamma^+(\alpha)|\psi\rangle = f(\alpha)$. Because $|\Gamma\rangle$ is a tight frame, this functional is $\frac{1}{\pi}\langle o(\alpha)|$, the coherent state function divided by the frame bound. That is not true for general frames, but they do have a pseudoinverse, and it defines a similar functional. In this general case, the kets $|\Gamma^+(\alpha)\rangle$ are generally known as the dual frame of the kets $|o(\alpha)\rangle$ [2].

We will consider various operators $|A\rangle$ that are related to $|\Gamma\rangle$, for example by restricting the integral to a finite volume of phase space. These will not be tight frames, and will not always span ket space. In these cases, where kets are being expanded approximately, we can define an error operator $E = (1 - |A\rangle\langle A^+|)$, so that when $f = \langle A^+|\psi\rangle$, the residual in the expansion is $E|\psi\rangle = |\psi\rangle - |A\rangle f$. Graphs of this operator can give hints as to how the expansion fails: for example, investigating how it expands over energy eigenvectors gives an idea of the energy cut off for the representation. However, the error operator does not give any information about the condition of the expansion. To find that out, we will usually have to look directly at the singular values of $|A\rangle$ and $|A^+\rangle$.

Figure 30 illustrates four expansion operators $|o(\alpha^1, \dots, \alpha^9)\rangle$, each of which is an expansion over 9 coherent states. The amplitudes α^1 through α^9 are marked by white circles on the complex plane. In the left column, these form a grid, in the right column, they are drawn at random from a square centred on the origin, with an area that generates the same average density as the grid to the left. The grid spacing is 0.7 in the top row, and 1.5 in the bottom row—this compares with the critical spacing of $\sqrt{\pi} = 1.77$, where a grid of coherent states is a basis for ket space [37].

The residual $\|(1 - |o(\alpha)\rangle\langle o^+(\alpha)|)|\beta\rangle\|^2$ when an arbitrary coherent state $|\beta\rangle$ is expanded over each ensemble is shown by grey shading at the point β . Coherent states whose amplitude is shaded black can be expanded exactly, those shaded white are orthogonal to every element of the frame, and those shaded grey have residuals intermediate between 0 and 1.

The red contours show the norm of the coefficient vector $\langle o(\alpha)^+|\beta\rangle$. This takes the value 1 at the amplitudes of the frame elements, where $\beta = \alpha^i$ and the expansion is the state itself. It could take smaller values, for example when $\beta = \alpha^i = \alpha^j$ and the smallest expansion is $|\beta\rangle = \frac{1}{2}(|\alpha^i\rangle + |\alpha^j\rangle)$ with $|c| = 1/\sqrt{2}$.

For each ensemble, the black region in which coherent states can be expanded exactly is a circle or square of area around 9. This is expected: the states in the ensemble are linearly independent, so they span a subspace of dimension 9, which will be close to the subspace spanned by the number states $|0\rangle$ through $|8\rangle$. In phase space, this subspace corresponds to the disk of radius 3 on which these kets have large components.

The dense sets of amplitudes lie in a square whose area is around 4.4, and the black area extends well outside it. The coherent states at the perimeter of the black area can still be expanded—their amplitudes are shaded black—but the red contours show that the expansion coefficients are large. This is necessary because the frame elements only contain small components of the number states around the perimeter of the black area. The norm $|c|$




Figure 30. Residuals (shading) and coefficient sizes (red contours) for expansion of a coherent state $|\beta\rangle$ over a set of 9 coherent states (amplitudes indicated by white circles).

rises to around 350 for the dense grid, and to a maximum of the same order for the random ensemble. The amplitudes in the sparse grid fill the black area of expandable states, so the elements include larger components of the number states around the perimeter. For this grid, the maximum of $|c|$ occurs at the perimeter, where it does for the dense grid, but it is not nearly as large.

Outside the perimeter, the expansion coefficients decrease. Coherent states with these large amplitudes are very nearly orthogonal to the ensemble components: their component in the span of $|\phi(\alpha)\rangle$ is a small ket that can be expanded with small coefficients.

The sparse random grid has some isolated amplitudes that happen to lie away from the others. In the region of phase space adjacent to one of these, say $\beta \approx \alpha^i$, the closest approximation to $|\beta\rangle$ is its projection parallel to $|\alpha^i\rangle$. There is only one nonzero expansion coefficient, with norm $|\langle \alpha^i | \beta \rangle|$. This forms a Gaussian blob of grey shading around the isolated amplitude.

41. Extrapolation in phase space

When frames were introduced in Section 12, the spaces were finite dimensional, and the expansion operator was a complex matrix. The sinc interpolants in Section 13 were an infinite set, but they were placed at discrete sample points, so that the set of coefficients

formed a discrete sequence.

The operator $|\Gamma\rangle$ represents kets by phase space distributions, functions defined on the whole complex plane. When the ket $|\psi\rangle$ represents an orbital occupied by many particles, the phase space representation with least norm, the distribution $\langle\Gamma^+|\psi\rangle$, is concentrated far from the origin. Such kets will occur when interacting particles are simulated, because the interaction energy is proportional to the square of the number of particles, so that the Hamiltonian operator amplifies many-particle components of the kets on which it acts. In order to represent kets on a computer, the distribution functions must be sampled on a finite region of the complex plane. If $\langle\Gamma^+|\psi\rangle$ extends beyond the sampled region, the ket can still be represented by another distribution, because $|\Gamma\rangle$ is not one-to-one. However, this distribution will take larger values and oscillate more sharply. Therefore it will have to be sampled carefully in order for the simulation to be numerically stable.

The inner product of coherent states with different amplitude,

$$\langle\alpha|\beta\rangle = e^{-\frac{1}{2}|\alpha|^2 - \frac{1}{2}|\beta|^2 + \beta^* \alpha},$$

first given in Equation 6, has a consequence for expansions. Suppose a coherent state $|\beta\rangle$ is approximated in the form in Equation 81, with variational parameters such that the error $\| |\beta\rangle - |o(\alpha)\rangle c \|$ is at most ϵ , but the amplitudes are placed with $|\beta - \alpha^r| > \delta$. Then $|\langle\beta|o(\alpha)\rangle c| > 1 - \epsilon$, by applying the Cauchy-Schwartz inequality to $|\beta\rangle$ and $|o(\alpha)\rangle c$. On the other hand $|\langle\beta|o(\alpha)\rangle c| = |\sum \langle\beta|\alpha^r\rangle c^r| \leq e^{-\frac{\delta^2}{2}} \sum |c_r| \leq \sqrt{R} e^{-\frac{\delta^2}{2}} \|c\|$. Therefore, as the amplitude β of the approximated state moves away from the variational amplitudes α , the coefficients $\|c\|$ must grow as $e^{\frac{\delta^2}{2}}$. The error condition requires that $\| |o(\alpha)\rangle c \| < 1 + \epsilon$, so the variational terms must almost entirely cancel. This bodes badly for the numerical stability of such expansions: it seems that states can be represented in a stable fashion only when the whole region of phase space that the state overlaps is covered by variational amplitudes. There remains the question of how densely it must be covered.

Although the operator $|\Gamma\rangle$ maps functions to kets, it still has a singular value expansion. The right singular vectors are functions, and the left singular vectors are kets, and the form of the expansion is

$$|\Gamma\rangle f = \sum_i |u_i\rangle \sigma_i \int v_i^* f,$$

where $\int g^* f$ is the L^2 inner product on complex functions. The singular values are real numbers. In fact, because the frame is tight, all the singular values are 1, we can choose any orthonormal basis to be the singular kets $|u_i\rangle$, and the corresponding singular functions are $v_i(\alpha) = \pi^{-\frac{1}{2}} \langle\alpha|u_i\rangle$. The singular functions do not span the space of square integrable functions, and $|\Gamma\rangle$ maps some functions to the zero ket.

Bargmann [34] noted that, for a fixed $|\psi\rangle = |\Gamma\rangle f$, the function f has the least L^2 norm when the function $e^{\frac{1}{2}|\alpha|^2} f(\alpha^*)$ is analytic. He deduced that the analytic functions form a Hilbert space under the inner product $\langle g, f \rangle = \int e^{-|\alpha|^2} g^*(\alpha) f(\alpha) d^2\alpha$. This Hilbert space is isomorphic to ket space, with $|\Gamma\rangle$ as the isometry. It is sometime convenient to work

with analytic functions, but for the discussion of frames we will stick with the expansion function f .

To think about extrapolation from a limited region of phase space [92], define a restricted operator

$$|\Gamma_\chi\rangle f = \pi^{-\frac{1}{2}} \int |\alpha\rangle \chi(\alpha) f(\alpha) d^2\alpha.$$

For example, χ could be a characteristic function, restricting the integral to its support. The operator $|\Gamma_\chi\rangle$ would then have some small singular values, where the support of the singular function v_i barely overlapped the support of χ .

Suppose a ket $|\psi\rangle$ was expanded over a region of phase space χ located distant from the origin, centred around an amplitude α . The expansion coefficients would be the same as those for the expansion of the ket $D^\dagger(\alpha)|\psi\rangle$ over the translation of this region to the origin. So it suffices to consider extrapolation from a disk around the origin of phase space. Luckily, in this case, the singular values turn out to values of a standard special function.

Suppose that χ is the characteristic function of a disk, with $\chi(\alpha) = 1$ when $|\alpha| \leq b$, and 0 otherwise. From symmetry, it can be guessed that the singular kets will be the number states $|n\rangle$ —this is verified numerically below. It can be shown by a Lagrange multiplier argument that the singular functions are the truncations to the disk of

$$f_n = \langle \Gamma^+ | n \rangle = \pi^{-\frac{1}{2}} \langle \alpha | n \rangle = \pi^{-\frac{1}{2}} e^{-\frac{1}{2}|\alpha|^2} \sum_{n=0}^{\infty} \frac{\alpha^{n*}}{\sqrt{n!}}.$$

The singular functions are scalar multiples of this that have unit norm when truncated. However, the singular values are ratios of norms, so the normalisation factors cancel out and can be ignored.

To evaluate the singular value σ_n , we will first evaluate $\|f_n\|$. We have

$$\begin{aligned} \|f_n\|^2 &= \frac{1}{\pi} \int_{|\alpha| \leq b} e^{-|\alpha|^2} \frac{|\alpha|^{2n}}{n!} d^2\alpha \\ &= \frac{1}{\pi} \int_0^b e^{-r^2} \frac{r^{2n}}{n!} 2\pi r dr \\ &= \frac{1}{n!} \int_0^{b^2} u^n e^{-u} du = \Gamma(n+1, b^2), \end{aligned}$$

where $\Gamma(n, x)$ is the incomplete gamma function. Everyone seems to define this function slightly differently, and the convention used here copies the one used in Matlab.

The norm of $|\Gamma_\chi\rangle f_n$ can be found from

$$\begin{aligned}
\| |\Gamma_\chi\rangle f_n \| |n\rangle &= |\Gamma_\chi\rangle f_n = \pi^{-\frac{1}{2}} \int_{|\alpha| \leq b} |\alpha\rangle f_n(\alpha) d^2\alpha \\
&= \frac{1}{\pi} \sum_{m=0}^{\infty} \int e^{-|\alpha|^2} \frac{\alpha^m}{\sqrt{m!}} |m\rangle \frac{\alpha^{n*}}{\sqrt{n!}} d^2\alpha \\
&= \frac{1}{\pi} \int_0^b r e^{-r^2} \sum_{m=0}^{\infty} \frac{r^{m+n}}{\sqrt{m!n!}} |m\rangle dr \int_0^{2\pi} e^{i(m-n)\phi} d\phi \\
&= \int_0^b 2r e^{-r^2} \frac{r^{2n}}{n!} |n\rangle dr \\
&= \Gamma(n+1, b^2) |n\rangle.
\end{aligned}$$

The singular value comes out to

$$\sigma_n = \frac{\| |\Gamma_\chi\rangle f_n \|}{\| f_n \|} = \sqrt{\Gamma(n+1, b^2)}. \quad (90)$$

Numerical experiments support the idea that the singular kets are number states, and the singular functions truncations of $\langle \alpha | n \rangle$. Figure 31 shows the singular functions for $|\Gamma_\chi\rangle$, where χ is the circle about the origin of radius $b = 5$. These were computed by taking the SVD of a matrix, where phase space was discretised on a square grid with pitch $h = 0.3$, and ket space was discretised by expanding over number states $|0\rangle$ through $|24\rangle$. This gives the first 25 singular kets and functions.

This demonstrates why such extrapolation is unstable. The function v_n that satisfies $|\Gamma_\chi\rangle v_n = \sigma_n |n\rangle$, and v_n has the smallest possible norm, is the restriction of $f(\alpha) = \langle \alpha | n \rangle$ to the disk χ . The function f takes its largest values on an annulus around the circle $|\alpha|^2 = n$. If n is small compared to b^2 , then all the coefficients fit inside the disk, $|\Gamma_\chi\rangle$ acts on v_n very nearly the same way as $|\Gamma\rangle$ does, and the singular value is 1 in both cases. For large n , only the inner edge of the annulus falls inside the set χ . This is still an expansion of $|n\rangle$, but it is an inefficient one with large coefficients, and the singular value is small.

The contour plot in Figure 32 shows how the analytic expression for the singular values σ_n decays with n during extrapolation from a disk of radius b . Reading along a horizontal line, we see that the largest singular values are close to one, then they drop to zero when n is around b^2 . This decay gets proportionately sharper as the disk gets larger. The rule of thumb is that there one large singular value for every unit circle of phase space, or area π .

The set of coherent states in any bounded region of phase space spans ket space. It is a frame for any set of kets that has bounded energy or particle number. However, as the bound on the energy increases, the lower frame bound shrinks, and some kets require expansion functions to take very large values. This will make the expansion operator ill-conditioned, and the differential equations derived from it stiff. If variational methods with coherent states are to be used numerically, the variational amplitudes must cover the region of phase space where the wave function has support.

This suggests that variational methods will not work for systems with many modes. In a high dimensional phase space, the volume of coherent states that overlap the wave function




Figure 31. Singular kets of the expansion operator over a disk in phase space.




Figure 32. Singular values of the expansion operator over a disk in phase space.

will be huge. For many types of state, it will increase exponentially with the dimension of the space. The next section considers the density at which phase space must be sampled; it turns out that, in order for the variational method to stable, the states must be as dense as if they formed a grid. Therefore the variational method it will not be practical for high dimensional systems. It might however be a useful method for systems of moderate dimension.

42. Density of coherent amplitudes

The last section discussed extrapolation from a bounded part of phase space to represent states in a larger part. It showed that, when kets are expanded over the coherent states within a bounded area, the expansion operator has one large singular value for each area π that the coherent states cover. This suggests that each such area corresponds to one independent ket; perhaps kets could be reconstructed exactly given one sample amplitude per area π , in a similar way that functions can be reconstructed exactly when they are sampled at their Nyquist density.

Another reason to suspect this might be possible is the angle between coherent states from Equation 7. Coherent states within a unit disk in quantum phase space point in similar directions in ket space to the state at its centre.

In fact, a discrete grid of coherent states with amplitudes $\alpha_{mn} = m\alpha_R + in\alpha_I$ does span phase space whenever the grid is dense enough that $\alpha_R\alpha_I \leq \pi$ [93, 37]. Some of these sources state the critical area as $x_0k_0 = 2\pi$, which is equivalent to an area π in quantum phase space given that $dx dk = 2d^2\alpha$. The critical density is shown in Figure 33, where each coherent state is drawn as a circle of area π , and the grid has spacing $\sqrt{\pi}$. It was later shown that at the critical density, where $\alpha_R\alpha_I = \pi$, the grid of coherent states is a basis for ket space, but there are states $|\psi\rangle$ for which the ratio $\|\langle\alpha_{mn}|\psi\rangle\|/\|\psi\rangle\|$ is arbitrarily small, and whose expansion coefficients must therefore be arbitrarily large [94]. So, in order for the grid of coherent states to be a frame, making expansions numerically stable, the density of the grid must be greater than the critical density.




Figure 33. A square grid of coherent states covers phase space.

Daubechies analysed frame bounds for grids of coherent states as they approached this limit [2]. The ratio of the frame bounds, which determines the condition of the expansion operator, increases quite sharply as the density approaches the π^{-1} limit. Also, the dual frame elements that determine the expansion coefficients approach singular functions. When phase space is densely sampled, and the frame is snug, the dual frame elements are very nearly the sampled coherent states themselves. So, if the sampled states cover the wave function, it is necessary to sample more densely than one state per unit circle of phase space.

As well as requiring a grid of coherent states to be be dense enough in phase space, Daubechies also found that the spacing of the grid points along the real axis must be similar

to that along the imaginary axis: the coherent states are circles in phase space, and the grid should be square rather than rectangular. Figure 34 illustrates why this is necessary. The grid is at the critical density, so the coherent states will be a basis for ket space. However, the momentum eigenstate shown as a dashed line has small components along all the coherent states in the grid, so its expansion coefficients will have to be large. This state is not normalisable, but a highly squeezed state exhibits the same problem. Daubechies did some numerical experiments, which suggest that the frame becomes significantly looser when the ratio of the real and imaginary grid spacings is around 2.




Figure 34. A rectangular grid of coherent states does not overlap all of phase space.

Daubechies derived some limits for the results for the frame bounds of a grid of coherent states formed from translations and momentum kicks of a wave packet that might not be Gaussian [2]. The main requirements for the frame to be snug are the the envelope $\sum_j |\psi(x - jx_0)|^2$ is close to a constant function of x , where $\psi_0(x)$ is the wave function for the vacuum state, and that a certain autocorrelation function is small, which ensures that the coherent states cover all momentum values.

Between the results of Section 41 and this section, the requirements for a discrete set of coherent states to form a snug frame are quite stringent. In order to avoid extrapolation, which requires large coefficients, the sample amplitudes must cover the disk of coherent states that can be expanded over them; this requires the amplitudes to be spread out so that each of them has an area π to itself. On the other hand, they can not be any more spread out than this, or else they stop spanning phase space.

The most basic result from Daubechies is that a tight frame of coherent states must be a grid in phase space. This is a very serious problem if the variation coherent state method is to be applied to cold atom systems. Simulations of cold atoms typically require dozens or hundreds of modes to be included in the simulated atomic field, resulting in a phase space with dozens or hundreds of dimensions. The number of points in a grid grows exponentially with the number of dimensions, and a phase space grid for an atomic field will be too large for practical computations.

The results here have considered regular grids of coherent states. For sinc wave packets, theorems exist about frame bounds of irregularly positioned wave packets [38]. Presumably similar theorems are possible for coherent states that are placed irregularly in phase space.

43. Discrete variational frame

The previous sections discussed expansions over coherent states. However, in the original variational problem, the derivative $-iH|\psi(z)\rangle$ is expanded over the frame $|\psi'(z)\rangle$. For the coherent state approximant of Equation 81,

$$|\psi'\rangle = (1 \quad a^\dagger) |\psi\rangle.$$

It follows that

$$\langle D_1\psi|D_1\psi\rangle = \langle\psi|\psi\rangle = e^{2\text{Re}\phi+|\alpha|}, \quad \langle D_2\psi|D_2\psi\rangle = (1+|\alpha|^2)\langle\psi|\psi\rangle,$$

and

$$\langle D_1\psi|D_2\psi\rangle = e^{2\text{Re}\phi+|\alpha|}\langle\alpha|a^\dagger|\alpha\rangle = \alpha^*\langle\psi|\psi\rangle.$$

The angle between $|D_1\psi\rangle$ and $|D_2\psi\rangle$ satisfies $\cot\theta = |\alpha|$. The frame comprises pairs of the partial derivative kets, which are orthogonal for the vacuum state, but approach parallel as $|\alpha|$ increases and $|\alpha|^2$ becomes proportionately close to $1+|\alpha|^2$. A pair of vectors like this was the first example in Section 12, and by the result derived there, the condition of $|\psi'\rangle = (|D_1\psi\rangle \quad |D_2\psi\rangle)$, with one coherent state, in the limit of large α , is approximately $|\alpha|^2$. The condition number of $|\psi'\rangle$ is at least as large as this, so the variational frame will loosen as it includes larger variational amplitudes. This will cause the variational equations of motion to become unnecessarily stiff.

Consider dynamical tunnelling, where the probability amplitude shifts from one region of phase space to another without passing through the points between. In terms of phase space, the values of $\langle\alpha|\Psi(t)\rangle$ change, but the region of phase space where this has support stays the same. If this process were expanded over a set of coherent states, then it might be expected that the variational amplitudes would stay constant, and the dynamics would be generated by changes in the variational coefficients. One consequence of the pairs of frame elements becoming parallel, with $|D_2\psi\rangle$ large compared to $|D_1\psi\rangle$, is that a change in a variational amplitude has a larger effect on $\langle\alpha|\psi(z)\rangle$ than a change of the same size in the variational coefficient, so that dynamic tunnelling might cause the variational amplitudes to shift instead of the coefficients.

This suggests that, although the least squares calculation might be more complicated if we parameterised the approximant as

$$|\psi(\alpha, c)\rangle = |o(\alpha)\rangle c,$$

the improvement in the condition of $|\psi'\rangle$ would make it worth doing.

The bounds for the variational frame will depend on the distribution of the coherent amplitudes and the size of the coefficients. Most obviously, if two amplitudes are nearly equal, the lower frame bound will be small. The two coefficients corresponding to these amplitudes can change in opposite ways, producing two Hilbert space vectors that cancel each other out, with very little effect on the expanded ket. If the coefficients corresponding to these amplitudes are already large and cancel, the upper bounds will become large

as well: a small change in one of the amplitudes, keeping everything else the same, will produce a large change in the expanded ket.

Figure 35 repeats Figure 30, except with an ensemble of 4 instead of 9, the density of the amplitude grid halved, and expansions made over the variational frame instead of the coherent states. The coefficients are smaller, because the states $a^\dagger|\alpha_r\rangle$ have components of higher number states than do $|\alpha_r\rangle$. The effect of randomness is larger, owing to the smaller ensemble size.




Figure 35. Residuals and coefficient sizes for variational frame expansions.

At this point, it should be clear that frames of coherent states, and variational frames composed from their partial derivatives, can very easily have loose bounds. The next sections will detail how this can lead to stiff equations of motion, and what it might be possible to do about it.

44. Regularisation

The Dirac-Frenkel equations of motion for the variational parameters $\zeta(t)$, for a quantum state whose dynamics are governed by a Hamiltonian H , have the form of Equation 87, $\zeta'(t) = \dot{\zeta}_H(\zeta(t))$. This equation of motion will be stiff whenever there is a way to perturb the parameters that significantly changes a component of the ket $H|\psi(\zeta)\rangle$ on the border of the region of phase space on which the variational frame $|\psi'(\zeta)\rangle$ is snug. In this case, it is not possible to compute the trajectory $\zeta(t)$ in a reasonable number of steps, so we

must look for a better conditioned problem that generates a less stiff equation of motion $w'(t) = \zeta_1(w(t))$. Provided that the state $|\psi(w(t))\rangle$ closely approximates $|\psi(\zeta(t))\rangle$, the variational parameters $w(t)$ that satisfy this equation can greatly differ from the original solution $\zeta(t)$. The variational problem is badly conditioned; there are ways to perturb $\zeta(t)$ that hardly change $|\psi'\rangle \cdot \zeta$, and functions such as ζ_1 do exist. The aim is to find one whose equation of motion is practical to solve.

This approach, where a badly conditioned problem is replaced by one that is physically equivalent but better conditioned, is known as regularisation. It is widely useful, and Section 45 will describe the contexts in which it is traditionally used. The remainder of this section will determine which parts of the variational problem need to be improved.

As explained in Section 11, the longest time step for which the usual integration formulae are stable is determined by the size of the Jacobian $\dot{\zeta}'_H$. The Dirac-Frenkel method sets

$$\dot{\zeta}'_H(z) = -i\langle\psi'(z)^+|H|\psi(z)\rangle,$$

so this will depend on the Jacobian of the mapping $f(z) = \langle\psi'(z)^+|$, a complicated second derivative of $|\psi\rangle$. We will suppose for simplicity of analysis that $\langle\psi'(z)^+|$ is independent of z , hoping that values of z that cause its derivative to be large also cause $\langle\psi'(z)^+|$ to be badly conditioned, so that the term in $\dot{\zeta}'_H$ proportional to its derivative is never much larger than the other terms. With this assumption, the function $\dot{\zeta}'_H$ is the composition of the nonlinear function $|\psi\rangle$ with the linear mappings H and $\langle\psi'(z)^+|$. The size of the Jacobian is largely determined by the norms of the linear mappings and of $|\psi'\rangle$.

The mapping $|\psi\rangle$ has a large derivative and is badly conditioned whenever the variational coefficients are large. In this case, as discussed in Section 43, small changes in the variational amplitudes α cause large changes in the state $|\psi(\alpha, \phi)\rangle$. In order for the coherent state ansatz to generate stable variational equations of motion, the size of the variational coefficients must remain modest. Section 15 showed that this condition, along with the requirement that the approximant $|\psi(\alpha, \phi)\rangle$ must stay close to the state $|\Psi\rangle$, requires the discrete variational amplitudes α to fill the volume of phase space where the overlap $\langle\alpha|\Psi\rangle$ is significant. The variational equations of motion will be reasonably stiff only when the coherent states cover phase space densely, and $|\phi(\alpha)\rangle$ forms a snug frame for the kets that can occur as $H|\psi\rangle$.

The linear operator H is its own derivative. Its size is a physical quantity, the maximum energy present in the state $|\psi(z)\rangle$, and it can not be reduced by numerical means without compromising the physical realism of the dynamics. However, the composition of H with $|\psi\rangle$ can have a worse condition than necessary when $|\psi'(z)\rangle$ is large and some of the singular kets belonging to its large singular values have large energies. When the variational coefficients are large and oscillating, so that the Jacobian $|\psi'\rangle$ is already large, the composition with H can make its condition worse.

The condition of $\langle\psi'(z)^+|$ is equal to the condition of $|\psi'(z)\rangle$, because their singular values are reciprocals. In the worst case, a singular ket belonging to a large singular value of $|\psi'(z)\rangle$ not only has a large energy, but is also rotated by H onto a singular ket belonging

to a small singular value of $|\psi'(z)\rangle$, which becomes a large singular value of $\langle\psi'(z)^+|$. The chance of this happening will presumably approach certainty as the number of large singular values increases.

Regularisation can not work miracles. In a high-dimensional phase space, the volume where $\langle\alpha|\Psi\rangle$ is significant will likely be too large for the discrete amplitudes α to fill it at the required density. Therefore any coefficients c that expand $|\Psi\rangle$ as $|o(\alpha)\rangle c$ will be large and oscillating, and so will any coefficients that closely approximate $|\Psi\rangle$. There is no hope of expanding an arbitrary state of many field modes as over a small set of coherent states: the only plausible way to represent such states is by fair sampling, where almost all of the parameters are recorded implicitly by assuming that they are similar to the few that are explicitly stored. A variational method with a fair sampling ansatz is an interesting idea, but outside the scope of this thesis.

We will conclude by discussing methods available today that might allow coherent states to be used as variational approximants in Heller's original context of moderate dimensional phase spaces. There are two related issues.

The first issue is to find a well conditioned replacement for $|\psi'(z)\rangle$, so that the $\zeta_1(z)$ which expands $H|\psi(z)\rangle$ is moderately sized and no more sensitive than necessary to small changes in $H|\psi(z)\rangle$. This type of issue, in which a badly-conditioned problem must be solved in isolation and the result will stand on its own, has been studied extensively and is well understood. There are standard methods to address it, which the following sections apply to variational dynamics.

The other issue is to ensure that $|\psi'(z)\rangle$ remains well conditioned as the simulation progresses. Variational dynamics does not solve one time step in isolation, rather it sets up a kind of feedback where the solution from one time step becomes the problem for the next one. The standard regularisation methods can avoid large and oscillating values of $\dot{\zeta}_1(\zeta(t))$, but this is not enough: each step $\Delta z = \dot{\zeta}_1(\zeta(t))\Delta t$ might be small, but there are a lot of steps, and they can add up to produce large and oscillating variational parameters $\zeta(t)$ at later times. These will cause $|\psi'(z)\rangle$ to become badly conditioned and the equations of motion to become stiff. So as well as regularising the least-squares problem at each time step, it is necessary to ensure that the variational parameters propagate in a way that keeps the problem well conditioned.

This type of issue has attracted much less interest than one-off regularisation. However, Fabčić's work offers a plausible approach. The function $|\psi\rangle$ becomes badly conditioned when the variational coefficients become large compared to the size of the quantum state they are representing; the obvious solution is to constrain their vector norm $|c|$ to not exceed a certain size limit. Implementing this is a step for the near future, and an solution will be sketched in Section 48.

45. Ill-posedness

Section 11 mentioned that differential operators have large singular values, because sinusoidal functions with large wave numbers are greatly magnified when they are differ-

entiated. This causes partial differential equations to reduce to stiff ordinary differential equations when they are discretised by a method of lines on a fine grid. It causes a more serious problem when a quantity is measured or sampled on a grid, and a derivative of the quantity is computed from the results—for example, we might want to know the speed of an object, but the only thing we can directly measure is its position at a set of discrete times.

The errors in the measurements vary randomly from one grid point to the next, so they include components with large wave numbers; when these components are magnified by differentiation, they can become larger than the derivative of the original quantity. Calculations in which this occurs are known as ill-posed problems, and they are the classic application of regularisation. The problem is illustrated schematically in Figure 36.

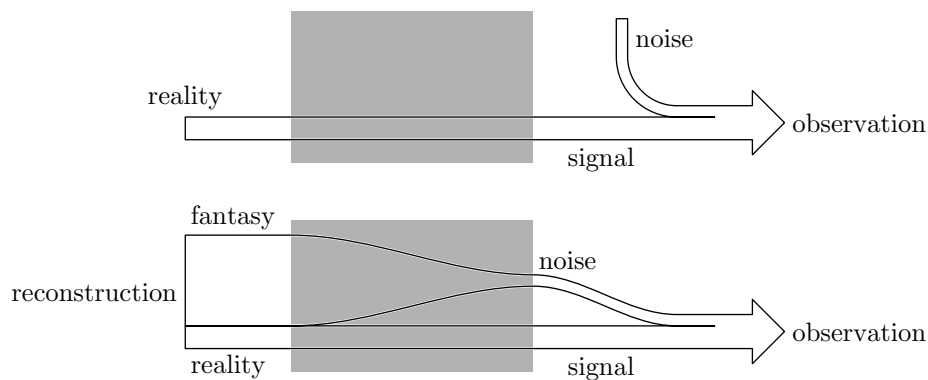


Figure 36. Schematic diagrams of signal and noise in an ill-posed problem.

The top figure shows the real relation between, for example, the speed of an object and the results of measurements of its position. There is a physical relationship between the speed and the position, represented by a grey process that turns reality into a signal. The process of measuring the signal combines it with some noise to produce the measurement results.

In contrast, the bottom figure shows the relation that is assumed when the results of position measurements are differentiated to determine a velocity. The derivative operator does not know that there are errors: it sees the measurement results, treats them as the true position of some physical object, and determines the velocity that this object would have required in order to pass exactly through the measured positions at the given times. The measured positions have Fourier components with large frequencies, due to the noise, so the derivative operator determines that such an object would have had to move very quickly.

This computed trajectory is not compatible with the law of inertia, and a smarter differentiation program could take this into account, allow for the errors, and return a realistic trajectory. This is called regularisation. Section 46 will discuss ways it can be done.

The effect of regularising the derivative operator is to make the computed derivative less sensitive to components of the function with large wave numbers. In the context of

partial differential equations, this is exactly what is required to improve the condition of the operator, and make the method of lines produce less stiff ordinary differential equations. In principle the stiff equations could be addressed by regularisation, although in practice it is much simpler and equally effective to compute the solution on a coarser grid, so that the components with large wavenumbers are discarded or aliased. However, when Schrödinger's equation is solved by a variational method, it is rarely possible to neatly characterise the components with large singular values, and regularisation becomes a practical approach.

The eigenvectors of differentiation are the complex exponentials, and the eigenvalues are their wavenumbers. Differentiating a function f with errors u can be expanded as

$$D(f + u)(x) = \frac{1}{2\pi} \int dk e^{ikx} k \int dx' e^{-ikx'} (f(x') + u(x')) = \frac{1}{2\pi} \int dk e^{ikx} k (F(k) + U(k)).$$

Assuming that the derivative Df exists, the coefficients $kF(k)$ must decay as the wave number increases, so that the integral $\int e^{ikx} kF(k)$ converged. However, there is no reason that the integral over $\int e^{ikx} kU(k)$ should converge. For instance, the integral does not converge when the errors u are due to white noise, making $U(k)$ a constant function; in this case, the derivative computed from the noisy data will diverge as $f(x)$ is sampled more frequently. The requirement that this does not happen, and the eigenvalues and noise components are such that their integral converges, is known as the Picard condition [95].

Differentiation has eigenvalues and eigenvectors, but in general there will be singular values in place of the k and singular vectors in place of the complex exponentials. In Section 15, we expanded the operator $|\Gamma_\chi\rangle$ from phase space functions to kets as a series of singular values and vectors

$$|\Gamma_\chi\rangle f = \sum_{i=1}^{\infty} |u_i\rangle \sigma_i \langle f, v_i\rangle.$$

Because $|\Gamma_\chi\rangle$ is a finite operator, this series will have to converge. The least squares problem was solved using the pseudoinverse, which can be expanded formally as

$$\langle \Gamma_\chi^+ | \psi \rangle = \sum_{i=1}^{\infty} v_i \sigma_i^{-1} \langle u_i | \psi \rangle.$$

It is not obvious that this series converges, because the restricted expansion operator $|\Gamma_\chi\rangle$ has many small singular values that tend to zero as i increases.

For the case where the set χ is a disk, the singular values given by Equation 90 decrease exponentially with i [96]. So it is possible to construct a normalisable state $|\psi\rangle = i^{-1} |u_i\rangle$, for which the coefficients $\sigma_i^{-1} \langle u_i | \psi \rangle$ increase exponentially as i increases. This certainly does not satisfy the Picard condition, and it is not surprising that these expansions are unstable.

The need for regularisation in variational dynamics differs from the traditional problems of deducing information from noisy measurements. The small singular values that occur in measurement problems have a physical origin: there exist sets of states that are

physically quite different, but that will produce nearly the same measurement data; when the details of implementation are stripped away, regularisation is based in a prejudice about which of these states is the most realistic one. In variational dynamics, on the other hand, the singular values are those of an expansion operator, and they are small because we have chosen to expand the states over a loose frame. However, there is no way to choose a tight frame for the whole physically accessible part of a high-dimensional phase space, at least not one that fits into the storage of a computer. We can try to constrain the variational frame to keep it as tight as possible, but it will inevitably have some small singular values in directions that overlap with the state being expanded, and the variational problem will have to be regularised.

This causes a difference in principle. The large and oscillating results that satisfy unregularised measurement problems are physically incorrect. They are attempts to find a physical state that would reproduce the errors in the experimental data. But those errors are produced by the process of measurement, not by the state being measured; the unregularised problem is unphysical, and the regularised problem is a better model of what is really occurring. In the variational problem, on the other hand, the large and oscillating solutions more nearly satisfy the laws of motion: the state that satisfies those laws has components with phase space amplitudes that are distant from the variational amplitudes, and the oscillating solutions represent those components more accurately than the regularised ones do. Regularising the variational problem is a compromise in which the original equations of motion are replaced by ones that are less realistic, but that numerical integration formulae are capable of solving.

46. Tikhonov conditioning

The type of ill-posed problem that occurs in when partial differential equations are discretised is known as the discrete type; the distinguishing feature is that the singular values of the ill-posed operator decrease continuously rather than suddenly dropping to zero. The standard way to regularise these problems is due to Tikhonov [97].

The problem

$$|\Gamma_\chi\rangle f = |\psi\rangle$$

can be regarded as a least squares problem, where $|\psi\rangle$ is chosen to minimise the residual $\| |\Gamma_\chi\rangle f - |\psi\rangle \|^2$. The exact solution has residual zero, and it is dominated by components belonging to the small singular values that make $|\Gamma_\chi\rangle$ badly conditioned and the variational equations stiff. Tikhonov modified this problem, choosing $|\psi\rangle$ to minimise the Tikhonov residual $\| |\Gamma_\chi\rangle f - |\psi\rangle \|^2 + \lambda \| f \|^2$. Here λ is a positive number, called a Lagrange parameter for reasons that will become clear shortly. Methods similar to this have been invented independently and applied to the Dirac-Frenkel problem [98].

There are two things to note. Firstly, the modified problem is still a linear least squares problem: it is equivalent to minimising

$$\left\| \begin{pmatrix} |\Gamma_\chi\rangle \\ \lambda I \end{pmatrix} f - \begin{pmatrix} |\psi\rangle \\ 0 \end{pmatrix} \right\|^2,$$

where I is the $2R \times 2R$ identity matrix. This can be solved by the usual least squares methods.

Secondly, suppose that f_λ is the solution for a given value of λ . If another function, g , had the same norm but a smaller residual, then f would not minimise the Tikhonov residual. Therefore Tikhonov regularisation is equivalent to constraining the norm of f , and seeking the solution with smallest residual given that norm. Or, equivalently, constraining the residual and seeking the solution with smallest norm. The parameter λ is a Lagrange parameter for this constrained optimisation problem.

Solving the Tikhonov problem for various values of λ generates a family of solutions. These can be plotted on a graph of residual against solution norm, where they mark the edge of the space of possible solutions. Above and to the right of the curve, a solution can be found with the given norm and residual: below and to the left, it can't. The boundary curve is known as an L-curve.

A typical L-curve that might result when measured data with errors are reduced using Tikhonov regularisation is shown in Figure 37. In the classic example, a derivative $g(x) = f'(x)$ is being computed from discrete samples $f_i = f(x_i) + e_i$, which include errors e_i . The top left is the limit where λ is large, constraining the norm of the solution $\|g\|$ to be zero, so that the regularised solution is $g = 0$, and the residual is the norm of the vector of samples f_i . As the constraint on the norm is relaxed, the solution becomes closer to the true value $g = f$, and the residual decreases. At the elbow of the curve, $\|g\|$ is constrained to be close to $\|f\|$, so the solution is close to the true value, and the residual is the norm of the measurement error $\|e\|$. Relaxing the constraint further allows spurious solutions that fit the measurement results f_i better than the true values $f(x_i)$ do. Typically, these won't fit much better, and the residual doesn't decrease much below $\|e\|$. However, the solution now includes components belonging to small singular values, and in the limit of small λ , at the lower right, the norm of the solution will increase as far as it is allowed to.




Figure 37. Regularisation L-curve for a normal ill-posed problem.

In a measurement problem, the elbow in the L-curve provides a way to choose the Lagrange parameter λ . When the regularised solution lies at the elbow on the curve, it contains all the components of the quantity that the data allow to be deduced: the remaining components belong to small singular values of the measurement process, and their effect on the results is smaller than the errors in the measurements. Fitting the measurement

results less closely would throw away information. On the other hand, fitting them more closely would not infer any more about the quantity being estimated, and would become an attempt to reproduce the measurement errors.

To demonstrate how this applies to expansions over coherent states, Figure 38 shows a cat state, and indicates the amplitudes of a set of 40 coherent states over which it can be expanded in the form $|\psi\rangle = |o(\alpha)\rangle \cdot c$. These amplitudes have been drawn from phase space, with the distribution $|\langle\alpha|\psi\rangle|^2$. Section 41 showed that the singular values of the restricted expansion operator $|\Gamma_\chi\rangle$ decay to very small values, whose singular vectors are number states that lie outside the sampled region χ . However, adding these singular vectors to numerical expansions with large coefficients can contribute to a better approximation of the expanded state.




Figure 38. Cat state sampling problem.

Therefore, as shown in Figure 39, the L-curve for this problem does not have an obvious elbow. As the regularisation is reduced, and the coefficient vector c is allowed to become larger, the approximation $|o(\alpha)\rangle \cdot c$ keeps converging towards $|\psi\rangle$. The L-curve does not contain enough information to choose a regularisation parameter; it suggests that less regularisation is always better. In variational dynamics, the constraint is the stiffness of the regularised equations of motion, which can not be inferred from the L-curve. This would need to be estimated directly, and sufficient regularisation chosen that the equations can be solved numerically. But ideally no more than that, so that the solution to the regularised equations satisfies the quantum law of motion as nearly as computationally possible.

Figure 40 shows a Picard plot for this problem. It displays three quantities: the singular values σ_r of the expansion operator $|o(\alpha)\rangle$, the components of the singular kets $|u_r\rangle$ in the expanded ket $|\psi\rangle$, and the components $\langle u_r|\psi\rangle/\sigma_r$ of the singular vectors v_i in the




Figure 39. L-curve for cat state sampling problem.

unregularised coefficient vector c . Note that, at the right of the graph, these components grow with r , despite the decay in the components $\langle u_r | \psi \rangle$. This suggests that regularisation is essential for this problem: as more amplitudes are added to the set α , the coefficient vector c that best approximates $|\psi\rangle$ will keep growing.

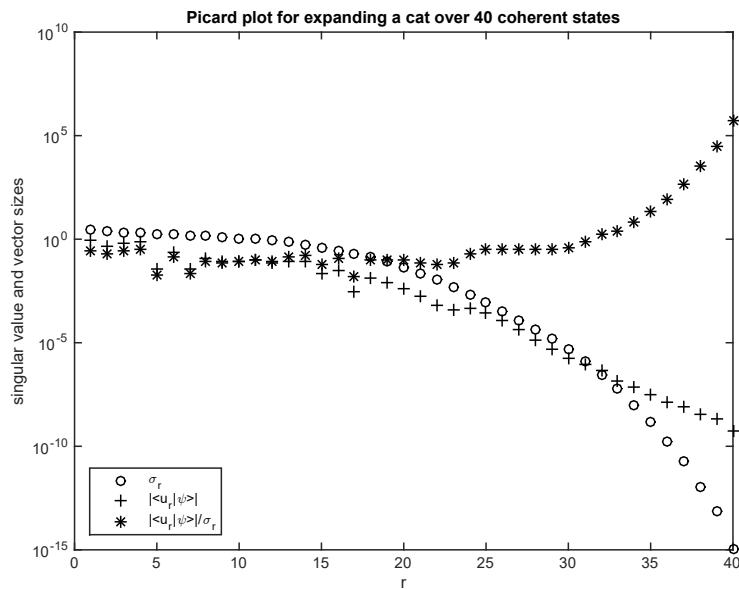


Figure 40. Picard plot for cat state sampling problem.

47. Effects of Tikhonov conditioning

As described above, the Tikhonov regularised variational problem assigns expansion coefficients d in order that

$$\begin{pmatrix} |\psi'(z)\rangle \\ \lambda I \end{pmatrix} d \approx \begin{pmatrix} -iH|\psi\rangle \\ 0 \end{pmatrix},$$

this equation being satisfied in a least-squares sense. The solution is given by

$$d = -i\langle\psi'(z)|_{\lambda}^{\dagger}|H|\psi\rangle = -i\begin{pmatrix} |\psi'(z)\rangle \\ \lambda I \end{pmatrix}^{\dagger} EH|\psi\rangle,$$

where \cdot^{\dagger} denotes the Moore-Penrose pseudoinverse of a matrix, $\cdot_{\lambda}^{\dagger}$ is the Tikhonov pseudoinverse, and E is the linear mapping

$$E|\psi\rangle = \begin{pmatrix} |\psi\rangle \\ 0 \end{pmatrix}.$$

Suppose that $|\psi'(z)\rangle$ has the singular-value decomposition $|\psi'(z)\rangle = |U\rangle SV^{\dagger}$, so that

$$\begin{pmatrix} |\psi'(z)\rangle \\ \lambda I \end{pmatrix} v_r = \begin{pmatrix} \sigma_r |u_r\rangle \\ \lambda v_r \end{pmatrix} = \sqrt{\sigma_r^2 + \lambda^2} w_r.$$

The orthonormality of the w_r follows from that of the $|u_r\rangle$ and v_r , so this provides a singular-value decomposition

$$\begin{pmatrix} |\psi'(z)\rangle \\ \lambda I \end{pmatrix} = WTV^{\dagger},$$

with singular values $\tau_r = \sqrt{\sigma_r^2 + \lambda^2}$. Therefore

$$\langle\psi'(z)|_{\lambda}^{\dagger} = \begin{pmatrix} |\psi'(z)\rangle \\ \lambda I \end{pmatrix}^{\dagger} E = VT^{-1}W^{\dagger}E,$$

and E selects the first column of W^{\dagger} , as

$$W^{\dagger}E = \begin{pmatrix} \frac{\sigma_1}{\sqrt{\sigma_1^2 + \lambda^2}} \langle u_1 | \\ \vdots \\ \frac{\sigma_R}{\sqrt{\sigma_R^2 + \lambda^2}} \langle u_R | \end{pmatrix}.$$

It follows that the right singular vectors of $\langle\psi'(z)|_{\lambda}^{\dagger}$ are the kets $|u_r\rangle$, and the singular values are

$$\sigma_r^{\dagger} = \frac{\sigma_r}{\tau_r \sqrt{\sigma_r^2 + \lambda^2}} = \frac{\sigma_r}{\sigma_r^2 + \lambda^2}.$$

This shows that Tikhonov conditioning divides the singular values of $|\psi'(z)\rangle$ into large ones and small ones, as compared to λ . For large σ_r , the singular value is very close to the unregularised inverse, with $\sigma_r^{\dagger} \approx 1/\sigma_r$. For small σ_r , the component of $|u_r\rangle$ is nearly discarded, with $\sigma_r^{\dagger} \approx \sigma_r/\lambda^2$. In this subspace, $|\psi'(z)|_{\lambda}^{\dagger} \approx \lambda^{-2}|\psi'(z)\rangle$. In between, there is a maximum at $\sigma_r = \lambda$, with $\sigma_r^{\dagger} = 1/2\lambda$, half the unregularised value.

This has consequences for the conditioning of the regularised problem, and how that varies with λ . When $\lambda < \sigma_R$, regularisation has very little effect: the condition of $|\psi'(z)^+|$ is the same as that of $|\psi'(z)\rangle$, which is often bad. As λ increases, so that $\sigma_1 < \lambda < \sigma_R$, the singular values lie on either side of the $\sigma_r = \lambda$ peak; the condition of $|\psi'(z)^+|$ is determined by how far below this peak σ_1 and σ_R fall, and the condition number has a sharp minimum when $\sigma_1^+ = \sigma_R^+$. When λ dominates the singular values, and $|\psi'(z)^+ \rangle = \lambda^{-2}|\psi'(z)\rangle$, it is again the case that their condition is the same. Of course, stiffness is determined by the size of the largest eigenvalues of the Jacobian, not the condition number.

This is illustrated in Figure 41, which plots the logarithms of the singular values of $|\psi'(z)\rangle$ against those of $|\psi'(z)^+|$, both scaled to the regularisation parameter λ . In the right hand plot, λ is small compared to the singular values, so these lie very close to the line $\sigma^+ = 1/\sigma$. In the left hand plot, λ is large, $|\psi'(z)^+ \rangle$ is nearly $\lambda^{-2}|\psi'(z)\rangle$, and the singular values lie on $\lambda\sigma^+ = \sigma/\lambda$. The singular values have been inverted, but this does not change the ratio of the smallest to the largest, so in both cases the condition numbers of $|\psi'(z)\rangle$ and $|\psi'(z)^+ \rangle$ are the same.




Figure 41. Effect of Tikhonov regularisation on singular values.

In the middle plot, λ is of the same order as the singular values, which is the usual case in well-regularised problems. In the illustrated case, where the singular values span several orders of magnitude, the condition number of $|\psi'(z)^+ \rangle$ is about the square root of that of $|\psi'(z)\rangle$. If the singular values were more even to begin with, and all fit in the curve at the top of the peak, the improvement would be greater. However, in this case it is unlikely that the problem would need to be regularised, so the square root improvement is likely the best that can be expected.

48. Conditioning with constraints

Fabčič [18] investigated a different way of avoiding stiffness in the variational equations of motion. While Tikhonov conditioning replaces the badly conditioned function $\dot{\zeta}_H(z)$ with a better conditioned one, this method constrains the variational parameters $\zeta(t)$ to values where partial derivatives $|D_i\psi\rangle$ form a snug frame, so that $\dot{\zeta}_H$ is already well conditioned around the values of z that actually occur as $\zeta(t)$.

This work did not systematically investigate the causes of the stiffness—it noted that the numerical integration routines sometimes required very short time steps, developed some heuristics about the sets of variational parameters for which this occurred, and constrained $\zeta(t)$ to avoid those kinds of parameters.

The main heuristic was that none of the variational coefficients was allowed to fall to a very small value. In this case, it is clear that $|\psi'(z)\rangle$ will be badly conditioned. When the coefficient of a coherent state is small, changes in its amplitude have a small effect on the superposition. However, this is not the only way that $|\psi'(z)\rangle$ can become nearly singular. For example, two coherent states with large coefficients and nearly equal amplitudes can cancel each other out. This occurs quite often when the quartic oscillator is simulated.

This approach partially succeeded. The stiffness occurred less often than with the unconstrained equations of motion, allowing some systems to be simulated that were unstable before. However, although stiffness occurred less often, it still occurred.

Knowing about frames and stiffness, it should be possible to take this approach, but to choose the constraints more systematically. The goal is to constrain z so that the variational frame $|\psi'(z)\rangle$ remains snug. The above analysis shows that this requires that the variational coefficients have reasonable magnitudes, that the variational amplitudes fill the region of phase space that overlaps the expanded state, and that they are sufficiently dense in this region. The constraints do not need to take all this into account. The reason that it is desirable for $|\psi'(z)\rangle$ to be a snug frame is simply that it allows the coefficients c in the expansion $|\Psi(\phi, \alpha)\rangle = |o(\alpha)\rangle \cdot c(\phi, \alpha)$ to have reasonable magnitudes. So it might be sufficient to constrain the norm of the coefficient vector $c(\phi, \alpha)$, and hope that the resulting constrained least squares problem forces the variational amplitudes to form a snug frame.

The mechanics of solving the constrained problem are detailed in [18]. Suppose that the constraint is $|c(\zeta(t))| \leq C$. At time steps when the parameters $\zeta(t)$ satisfy the constraint with $|c(\zeta(t))| < C$, the variational method can proceed unconstrained. It will probably require some Tikhonov type regularisation to avoid stiffness due to small variational coefficients, and also to limit the step $\Delta\zeta$ in one time step to a size where the differential approximation $|\psi(\zeta + \Delta\zeta)\rangle \approx |\psi(\zeta)\rangle + |\psi'(\zeta)\rangle \cdot \Delta\zeta$ is valid. At other time steps, when the constraint is only just satisfied and $|c(\zeta(t))| = C$, there are two possibilities. In the case that the unconstrained $\dot{\zeta}_H(t)$ points in a direction that reduces $|c(\zeta(t))|$, there is no problem. Otherwise, it is necessary to find a $\dot{\zeta}_H(t)$ that satisfies the variational problem as well as possible, while keeping $|c(\zeta(t))|$ constant. In other words, we need $\dot{\zeta}_H(t)$ to be orthogonal to the gradient of c . Although c is not an analytic function, this gradient can be expressed either in terms of real and imaginary parts, or in terms of Wirtinger derivatives [99].

49. Other approximants

Besides the expansions over coherent states that we have been studying, there is another class of approximants in common use for simulations of quantum chemistry and atom optics. This is paradoxically named the multiconfigurational Hartree approach [86]. It was proposed in 1990, since when it has been used quite widely, and with considerable success in quantum chemistry applications with moderate numbers of particles [100].

The approximant is usually presented as an expansion of a wave function in first quantisation with many coordinates; in our context, it would expand a wave function in second quantisation over many modes, but the form is identical. The ansatz comprises an orthonormal set of wave functions for each mode, $\phi_1^{(i)}$ through $\phi_{r_i}^{(i)}$, where i indexes the modes. The approximant combines products of these functions with arbitrary coefficients, as

$$\psi(x_1, x_2, \dots, x_n) = \sum_{j=1}^{r_1} \cdots \sum_{l=1}^{r_n} a_{j\dots l} \phi_j^{(1)}(x_1) \cdots \phi_l^{(n)}(x_n).$$

Numerically, the functions ϕ are stored as sets of sample values, or expanded over some fixed basis for each mode.

When this form has been used numerically, it has been constrained in order to avoid the kind of stiffness that occurs in coherent state expansions. The set of functions $\phi_j^{(i)}$ for a fixed mode j are usually required to be orthonormal. Also, when the Dirac-Frenkel method is applied, the rate of change $\partial_t \phi_j^{(i)}$ is constrained to be orthogonal to all the $\phi_j^{(i)}$. This does not lose any generality—any component parallel to these functions can be absorbed into their coefficients—but it reduces the redundancy in the representation and improves the condition of the variational frame. Some regularisation still had to be applied. For example, when the expanded state actually is a Hartree state, all but one of the coefficients should be zero, say $a_{11\dots 1}$. Then the dynamics of $\phi_i^{(m)}$ with $i > 1$ are not determined by the variational method.

The Gaussian wave packet approximant is a special case of this, where all coefficients are zero except for the diagonal $a_{jj\dots j}$, and the wave function $\phi_j^{(1)}(x_1) \cdots \phi_j^{(n)}(x_n)$ is a coherent wave packet $|\alpha_j^{(1)}\rangle \otimes \cdots \otimes |\alpha_j^{(n)}\rangle$. For practical purposes, the question is how rapidly the approximant converges to the exact state as the number of parameters increases; whether the coherent state or the multiconfigurational form converges more rapidly will no doubt depend on the state to be represented. When the state is actually a Hartree state, for example, the multiconfigurational form is as efficient as could reasonably be expected, but the coherent state form will require a very large grid of amplitudes covering the high dimensional phase space. Unfortunately, the orthogonality constraints that stabilise the multiconfigurational approximant numerically are not compatible with the ϕ_j continuing to have the form of coherent states.

There are elaborations of this form. Most of these apply it recursively, with functions $\phi_i^{(j)}(x_1, \dots, x_n)$ that are themselves expanded in multiconfigurational form [101]. This has produced converged solutions to the spin boson problem [102, 103], so it has some potential in quantum field applications.

However, other applications of this method to quantum fields have been less successful. As more configurations are added, this method converges to an exact representation of the quantum state, and the amount of computer storage required to exactly represent a quantum field increases exponentially with the number of particles and modes. In practice, only a handful of modes can be treated, and the convergence can not be tested. Some published results have turned out to be wrong due to an inadequate number of modes being used [104, 105, 81].

The truncated Wigner simulations of one-dimensional cold atom experiments performed as expected. In one set of simulations, a coherent state was sampled with a constant order parameter as the field amplitude. This was not a steady state under truncated Wigner dynamics. The coherence function $g^{(2)}$ relaxed from 1, the coherent state value, towards the Bose field ground state value. Meanwhile, the density of bosons in momentum space oscillated, and the dispersion relation was consistent with these oscillations being caused by Bogoliubov sound waves.

In a second set of simulations, samples were drawn from a Bogoliubov approximate ground state instead of a coherent one. The oscillations in the momentum density were much smaller, suggesting that this is a better approximation to the ground state. For values of γ above 0.1, there was still some noticeable relaxation of the coherence function, and the oscillations appeared to follow the Bogoliubov dispersion curve. This is surprising, because the Bogoliubov sound waves should not be excited in this state. The Bogoliubov state was calculated using the density of the condensate mode, without correcting for the repulsion from non-condensate bosons. For larger values of γ , the non-condensate fraction becomes appreciable, and it is possible that the initial state was simply calculated with the wrong density. The Bogoliubov state can be modified to account for non-condensate bosons [106], and this might reduce the oscillations further.

There are some alternative approaches to drawing Wigner samples for the ground state of the atomic field. One of them involves sampling the field in a coherent state, with the repulsion of the atoms initially set to zero, so that the coherent state is the exact ground state of the atomic field. The atoms are then simulated as the repulsion increases adiabatically [107]. This method was reported to be more effective in open systems, where atoms and energy were exchanged between the field and some thermal reservoir, than for closed systems such as the split traps.

The other method is a somewhat similar approach [108]. The field is sampled in an arbitrary state—a coherent state is convenient—and a simulation is carried out of the dynamics of this field in contact with a thermal reservoir. In principle, such a simulated field should approach thermal equilibrium, at which point it will generate an ensemble of samples for the thermal state of the field. At low temperature, this will be the ground state. This method has been derived on paper, and numerical simulations have been attempted, but these have not yet been made to converge.

The sensible approach to truncated Wigner simulations of these tightly-confined gasses would be to develop all three methods in parallel. It is possible that the convergence issues in each method have some common cause, or possibly one method might succeed.

This goal of this project was to simulate the dynamics of atomic fields, by expanding the quantum state of the field as a superposition of multimode coherent states. Chapter 5 argued that this can not be done. The multimode coherent amplitudes lie in a high dimensional phase space. There is a special case where the state of the field is exactly coherent, or where the coherent amplitudes of the modes are very strongly correlated, so that the state can be expanded over a few coherent states. In general, the coherent amplitudes for which the state has significant components will fill a volume of phase space, and the number of coherent states required to cover this volume will increase exponentially with the number of modes. It was shown in Section 4.1 that the coherent amplitudes must fill the volume of phase space overlapping the expanded state, or else the expansion will be numerically unstable, and the variational equations of motion will be impractically stiff. So the variational method, in the general case, is no more efficient than a grid of coherent amplitudes in phase space.

The variational method is still worth pursuing in its original context, systems with a dozen or so degrees of freedom. It has not been used as widely as it might, because attempts to employ it keep failing due to numerical instability. Hopefully, the analysis that demonstrates why the method can not simulate quantum fields will allow these problems to be addressed more systematically, and the method might be made reliable enough to be used more widely in the future.

References

- [1] C. J. Pethick and H. Smith. *Bose-Einstein Condensation in Dilute Gases*. Cambridge University Press (Cambridge), 2002.
- [2] Ingrid Daubechies. *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics, 1992.
- [3] Lloyd N. Trefethen. *Spectral Methods in Matlab*. Society for Industrial and Applied Mathematics, 2000.
- [4] Alexander L. Fetter and John Dirk Walecka. *Quantum Theory of Many-Particle Systems*. McGraw-Hill (New York), 1971.
- [5] Elliott H. Lieb and Werner Liniger. Exact analysis of an interacting Bose gas. I. the general solution and the ground state. *Physical Review*, 130:1605, 1963.
- [6] C. N. Yang and C. P. Yang. Thermodynamics of a one-dimensional system of bosons with repulsive delta-function interaction. *Journal of Mathematical Physics*, 10:1115, 1969.
- [7] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven. Regimes of quantum degeneracy in trapped 1D gases. *Physical Review Letters*, 85:3745, 2000.
- [8] Takuya Kitagawa, Adilet Imambekov, Jörg Schmiedmayer, and Eugene Demler. The dynamics and prethermalization of one-dimensional quantum systems probed through the full distributions of quantum noise. *New Journal of Physics*, 13:073018, 2011.
- [9] J. Denschlag, J. E. Simsarian, D. L. Feder, Charles W. Clark, L. A. Collins, J. Cubizolles, L. Deng, E. W. Hagley, K. Helmerson, W. P. Reinhardt, S. L. Rolston, B. I. Schneider, and W. D. Phillips. Generating solitons by phase engineering of a Bose-Einstein condensate. *Science*, 287:97, 2000.
- [10] T. P. Simula, P. Engels, I. Coddington, V. Schweikhard, E. A. Cornell, and R. J. Ballagh. Observations on sound propagation in rapidly rotating Bose-Einstein condensates. *Physical Review Letters*, 94:080404, 2005.

- [11] U. Schollwöck. The density-matrix renormalization group. *Reports on Mathematical Physics*, 77:259, 2005.
- [12] P. B. Blakie, A. S. Bradley, M. J. Davis, R. J. Ballagh, and C. W. Gardiner. Dynamics and statistical mechanics of ultra-cold Bose gases using c-field techniques. *Advances in Physics*, 57:363, 2008.
- [13] E. J. Heller. Time-dependent approach to semiclassical dynamics. *Journal of Chemical Physics*, 62:1544, 1975.
- [14] Eric J. Heller. Time dependent variational approach to semiclassical dynamics. *Journal of Chemical Physics*, 64:63, 1976.
- [15] Robert Heather and Horia Metiu. A numerical study of the multiple Gaussian representation of time dependent wave functions of a Morse oscillator. *Journal of Chemical Physics*, 84:3250, 1986.
- [16] Rex T. Skodje and Donald G. Truhlar. Localized Gaussian wave packet methods for inelastic collisions involving anharmonic oscillators. *Journal of Chemical Physics*, 80:3123, 1984.
- [17] Eric J. Heller. Frozen Gaussians: A very simple semiclassical approximation. *Journal of Chemical Physics*, 75:2923, 1981.
- [18] Tomaž Fabčić, Jörg Main, and Günter Wunner. Time propagation of constrained coupled Gaussian wave packets. *Journal of Chemical Physics*, 128:044116, 2008.
- [19] Julian Grond, Gregory von Winckel, Jörg Schmiedmayer, and Ulrich Hohenester. Optimal control of number squeezing in trapped Bose-Einstein condensates. *Physical Review A*, 80:053625, 2009.
- [20] William Henry Louisell. *Quantum Statistical Properties of Radiation*. John Wiley and Sons, 1973.
- [21] Erwin Schrödinger. Der stetige übergang von der mikro-zur makromechanik. *Naturwissenschaften*, 14, 1926.
- [22] Roy J. Glauber. Coherent and incoherent states of the radiation field. *Physical Review*, 131:2766, 1963.
- [23] E. C. G. Sudarshan. Equivalence of semiclassical and quantum mechanical descriptions of statistical light beams. *Physical Review Letters*, 10:277, 1963.
- [24] J. E. Moyal. Quantum mechanics as a statistical theory. *Mathematical Proceedings of the Cambridge Philosophical Society*, 45:99, 1949.
- [25] Laura E. C. Rosales-Zárate and P. D. Drummond. Probabilistic Q-function distributions in fermionic phase-space. *New Journal of Physics*, 17:032002, 2015.
- [26] Paul Aiden Morris Dirac. *The Principles of Quantum Mechanics*. Oxford University Press, 1930.

- [27] Charles Kittel and Herbert Kroemer. *Thermal Physics*. W. H. Freeman, 1980.
- [28] M. Gring, M. Kuhnert, T. Langen, T. Kitagawa, B. Rauer, M. Schreitl, I. Mazets, D. Adu Smith, E. Demler, and J. Schmiedmayer. Relaxation and prethermalization in an isolated quantum system. *Science*, 337:1318, 2012.
- [29] Waseem S. Bakr, Amy Peng, M. Eric Tai, Ruichao Ma, Jonathon Simon, Jonathon I. Gillen, Simon Fölling, Lode Pollet, and Markus Greiner. Probing the superfluid to Mott insulator transition at the single atom level. *Science*, 329:547, 2010.
- [30] Waseem S. Bakr, Jonathon I. Gillen, Amy Peng, Simon Fölling, and Markus Greiner. A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice. *Nature*, 462:74, 2009.
- [31] Clément Sayrin, Igor Dotsenko, Xingxing Zhou, Bruno Peaudecerf, Théo Rybarczyk, Sébastien Gleyzes, Pierre Rouchon, Mazyar Mirrahimi, Hadis Amini, Michel Brune, Jean-Michel Raimond, and Serge Haroche. Real-time quantum feedback prepares and stabilizes photon number states. *Nature*, 477:73, 2011.
- [32] Jutho Haegeman, Damian Draxler, Vid Stojevic, J. Ignacio Cirac, Tobias J. Osborne, and Frank Verstraete. Quantum Gross-Pitaevskii equation. arXiv preprint number 1501.06575, 2016.
- [33] Elias Wegert. *Visual Complex Functions: An Introduction with Phase Portraits*. Birkhäuser/Springer Basel AG, 2012.
- [34] V. Bargmann. On a Hilbert space of analytic functions and an associated integral transform. *Communications on Pure and Applied Mathematics*, 14:187, 1961.
- [35] Lloyd N. Trefethen and David Bau III. *Numerical Linear Algebra*. SIAM, 1997.
- [36] Lloyd N. Trefethen. Finite difference and spectral methods for ordinary and partial differential equations. available at <http://people.maths.ox.ac.uk/trefethen/pdetext.html>, 1996.
- [37] John von Neumann. *Mathematical foundations of quantum mechanics*. Princeton University Press, 1955.
- [38] R. J. Duffin and A. C. Schaeffer. A class of nonharmonic fourier series. *Transactions of the American Mathematical Society*, 72:341, 1952.
- [39] K. E. Cahill and Roy J. Glauber. Ordered expansions in boson amplitude operators. *Physical Review*, 177:1857, 1968.
- [40] K. E. Cahill and R. J. Glauber. Density operators and quasiprobability distributions. *Physical Review*, 177:1882, 1968.
- [41] Freeman Dyson. *Advanced Quantum Mechanics*. World Scientific (Singapore), 2011.
- [42] Howard Carmichael. *Statistical Methods in Quantum Optics 1: Master Equations and Fokker-Planck Equations*. Springer-Verlag (Berlin), 1999.

- [43] Mark R. Dowling, Peter D. Drummond, Matthew J. Davis, and Piotr Deuar. Time-reversal test for stochastic quantum dynamics. *Physical Review Letters*, 94:130401, 2005.
- [44] E. Wigner. On the quantum correction for thermodynamic equilibrium. *Physical Review*, 40:749, 1932.
- [45] Eric J. Heller. Wigner phase space method: Analysis for semiclassical applications. *Journal of Chemical Physics*, 65:1289, 1976.
- [46] P. D. Drummond and P. Kinsler. Quantum tunneling and thermal activation in the parametric oscillator. *Physical Review A*, 40:4813, 1989.
- [47] R. Graham and F. Haake. *Quantum Statistics in Optics and Solid-State Physics*. Springer, 1973.
- [48] Peter D. Drummond and A. D. Hardman. Simulation of quantum effects in Raman-active waveguides. *Europhysics Letters*, 21:279, 1993.
- [49] M. J. Steel, M. K. Olsen, L. I. Plimak, P. D. Drummond, S. M. Tan, M. J. Collett, D. F. Walls, and R. Graham. Dynamical quantum noise in trapped Bose-Einstein condensates. *Physical Review A*, 58:4824, 1998.
- [50] Crispin W. Gardiner and Peter Zoller. *Quantum Noise*. Springer-Verlag (Berlin), 2000.
- [51] Crispin Gardiner. *Stochastic Methods: a Handbook for the Natural and Social Sciences*. Springer (Berlin), 2009.
- [52] P. D. Drummond and C. W. Gardiner. Generalised P-representations in quantum optics. *Journal of Physics A*, 13:2353, 1980.
- [53] A. Gilchrist, C. W. Gardiner, and P. D. Drummond. Positive P representation: Application and validity. *Physical Review A*, 55:3014, 1997.
- [54] J. F. Corney and P. D. Drummond. Gaussian phase-space representations for fermions. *Physical Review B*, 73:125112, 2006.
- [55] Takeshi Aimi and Masatoshi Imada. Does simple two-dimensional Hubbard model account for high-Tc superconductivity in copper oxides? *Journal of the Physical Society of Japan*, 76:113708, 2007.
- [56] L. I. Plimak, M. K. Olsen, M. Fleischhauer, and M. J. Collett. Beyond the Fokker-Planck equation: Stochastic simulation of complete Wigner representation for the optical parametric oscillator. *Europhysics Letters*, 56:372, 2001.
- [57] Shin-Ichi Sawada, Robert Heather, Bret Jackson, and Horia Metiu. A strategy for time dependent quantum mechanical calculations using a Gaussian wave packet representation of the wave function. *Journal of Chemical Physics*, 83:3009, 1985.

- [58] C. A. Sackett, J. M. Gerton, M. Welling, and R. G. Hulet. Measurements of collective collapse in a Bose-Einstein condensate with attractive interactions. *Physical Review Letters*, 82:876, 1999.
- [59] Alexander L. Fetter and Anatoly A. Svidzinsky. Vortices in a trapped dilute Bose-Einstein condensate. *Journal of Physics: Condensed Matter*, 13:R135, 2001.
- [60] R. Lopes, A. Imanaliev, A. Aspect, M. Cheneau, D. Boiron, and C. I. Westbrook. Atomic Hong-Ou-Mandel experiment. *Nature*, 520:66, 2015.
- [61] N. J. van Druten and Wolfgang Ketterle. Two-step condensation of the ideal Bose gas in highly anisotropic traps. *Physical Review Letters*, 79:549, 1997.
- [62] A. Görlitz, J. M. Vogels, A. E. Leanhardt, C. Raman, T. L. Gustavson, J. R. Abo-Shaer, A. P. Chikkatur, S. Gupta, S. Inouye, T. Rosenband, and W. Ketterle. Realization of Bose-Einstein condensates in lower dimensions. *Physical Review Letters*, 87:130402, 2001.
- [63] V. N. Popov. *Functional Integrals in Quantum Field Theory and Statistical Physics*. D. Reidel (Dordrecht), 1983.
- [64] Rajibul Islam, Ruichao Ma, Philipp M. Preiss, M. Eric Tai, Alexander Lukin, Matthew Rispoli, and Markus Greiner. Measuring entanglement entropy in a quantum many-body system. *Nature*, 528:77, 2015.
- [65] J. Reichel and V. Vuletić. *Atom Chips*. Wiley-VCH (Weinheim), 2011.
- [66] T. Schumm, S. Hofferberth, L. M. Andersson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Krüger. Matter-wave interferometry in a double well on an atom chip. *Nature Physics*, 1:57, 2005.
- [67] D. Pines. *The Many-Body Problem*. W. A. Benjamin, 1961.
- [68] I. Lesanovsky, T. Schumm, S. Hofferberth, L. M. Andersson, P. Krüger, and J. Schmiedmayer. Adiabatic radio frequency potentials for the coherent manipulation of matter waves. *Physical Review A*, 73:033619, 2006.
- [69] Karen. V. Kheruntsyan, D. M. Gangardt, Peter D. Drummond, and G. V. Shlyapnikov. Pair correlations in a finite-temperature 1D Bose gas. *Physical Review Letters*, 91:040403, 2003.
- [70] T. Holstein and H. Primakoff. Field dependence of the intrinsic domain magnetization of a ferromagnet. *Physical Review*, 58:1098, 1940.
- [71] Lorenzo Isella and Janne Ruostekoski. Nonadiabatic dynamics of a Bose-Einstein condensate in an optical lattice. *Physical Review A*, 72:011601, 2005.
- [72] D. A. W. Hutchinson, E. Zaremba, and A. Griffin. Finite temperature excitations of a trapped Bose gas. *Physical Review Letters*, 78:1842, 1997.

- [73] Alexander L. Fetter. Nonuniform states of an imperfect Bose gas. *Annals of Physics*, 70:67, 1972.
- [74] Jean-Sébastien Caux, Pasquale Calabrese, and Nikita A Slavnov. One-particle dynamical correlations in the one-dimensional Bose gas. *Journal of Statistical Mechanics: Theory and Experiment*, 2007:P01008, 2007.
- [75] J. Ruostekoski and L. Isella. Dissipative quantum dynamics of bosonic atoms in a shallow 1D optical lattice. *Physical Review Letters*, 95:110403, 2005.
- [76] Tomasz Świsłocki and Piotr Deuar. Quantum fluctuation effects on the quench dynamics of thermal quasicondensates. *Journal of Physics B*, 49:145303, 2016.
- [77] XSPDE software. World wide web page accessed 7th February 2017, at https://github.com/peterddrummond/xspde_matlab.
- [78] Simon Kieseewetter, Rodney Polkinghorne, Bogdan Opanchuk, and Peter D. Drummond. xspde: Extensible software for stochastic equations. *SoftwareX*, 5:12, 2016.
- [79] Greg Collecutt and Peter D. Drummond. xmds: eXtensible multi-dimensional simulator. *Computer Physics Communications*, 142:219, 2001.
- [80] Graham R. Dennis, Joseph J. Hope, and Mattias T. Johnsson. XMDS2: Fast, scalable simulation of coupled stochastic partial differential equations. *Computer Physics Communications*, 184:201208, 2013.
- [81] Peter D. Drummond and Bogdan Opanchuk. Truncated Wigner dynamics and conservation laws. *Physical Review A*, 96:043616, 2017.
- [82] D. R. Hartree and W. Hartree. Self-consistent field, with exchange, for Beryllium. *Proceedings of the Royal Society*, 150:9, 1935.
- [83] Takuya Kitagawa, Susanne Pielawa, Adilet Imambekov, Jörg Schmiedmayer, Vladimir Gritsev, and Eugene Demler. Ramsey interference in one-dimensional systems: The full distribution function of fringe contrast as a probe of many-body dynamics. *Physical Review Letters*, 104:255302, 2010.
- [84] J. M. Luttinger. An exactly soluble model of a many-fermion system. *Journal of Mathematical Physics*, 4:1154, 1963.
- [85] Daniel C. Mattis and Elliott H. Lieb. Exact solution of a many-fermion system and its associated boson field. *Journal of Mathematical Physics*, 6:304, 1965.
- [86] H.-D. Meyer, U. Manthe, and L. S. Cederbaum. The multi-configurational time-dependent Hartree approach. *Chemical Physics Letters*, 165:73, 1990.
- [87] Bogdan Opanchuk, Rodney Polkinghorne, Oleksandr Fialko, Joachim Brand, and Peter D. Drummond. Quantum simulations of the early universe. *Annalen der Physik*, 525:866, 2013.

- [88] B. M. Caradoc-Davies. *Vortex Dynamics in Bose-Einstein Condensates*. PhD thesis, University of Otago, 2000.
- [89] Michael Spivak. *Calculus on Manifolds*. Westview Press, 1965.
- [90] J. Frenkel. *Wave Mechanics*. Clarendon (Oxford), 1934.
- [91] A.D. McLachlan. A variational solution of the time-dependent Schrodinger equation. *Molecular Physics*, 8:39, 1964.
- [92] Ingrid Daubechies. Time-frequency localization operators: a geometric phase space approach. *IEEE Transactions on Information Theory*, 34:605, 1988.
- [93] V. Bargmann, P. Butera, L. Girardello, and John R. Klauder. On the completeness of the coherent states. *Reports on Mathematical Physics*, 2:221, 1971.
- [94] H. Bacry, A. Grossman, and J. Zak. Proof of the completeness of lattice states in the kq-representation. *Physical Review B*, 12:1118, 1975.
- [95] Per Christian Hansen. *Discrete inverse problems: insight and algorithms*. Society for Industrial and Applied Mathematics, 2010.
- [96] The Wolfram functions site. World wide web page accessed 7th February 2017, at <http://functions.wolfram.com/06.06.06.0042.01>.
- [97] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-Posed Problems*. Winston (Washington D.C.), 1977.
- [98] Kenneth G. Kay. The matrix singularity problem in the time-dependent variational method. *Chemical Physics*, 137:165, 1989.
- [99] Are Hjørungnes and David Gesbert. Complex-valued matrix differentiation: Techniques and key results. *IEEE Transactions on Signal Processing*, 55:2740, 2007.
- [100] H.-D. Meyer, G. A. Worth, and F. Gatti. *Multidimensional Quantum Dynamics: MCTDH Theory and Applications*. Wiley-VCH, 2009.
- [101] Haobin Wang and Michael Thoss. Multilayer formulation of the multiconfiguration time-dependent Hartree theory. *Journal of Chemical Physics*, 119:1289, 2003.
- [102] Haobin Wang. Basis set approach to the quantum dissipative dynamics: Application of the multiconfiguration time-dependent Hartree method to the spin-boson problem. *Journal of Chemical Physics*, 113:9948, 2000.
- [103] Haobin Wang and Michael Thoss. From coherent motion to localization: dynamics of the spin-boson model at zero temperature. *New Journal of Physics*, 10:115005, 2008.
- [104] J. G. Cosme, C. Weiss, and J. Brand. Center-of-mass motion as a sensitive convergence test for variational multimode quantum dynamics. *Physical Review A*, 94:043603, 2016.

- [105] B. Opanchuk and P. D. Drummond. One-dimensional bose gas dynamics: breather relaxation. *Physical Review A*, 96:053628, 2007.
- [106] R. J. Dodd, Mark Edwards, Charles W. Clark, and K. Burnett. Collective excitations of bose-einstein-condensed gases at finite temperatures. *Physical Review A*, 57:r32, 1998.
- [107] Janne Ruostekoski and A. D. Martin. Truncated Wigner method for Bose gasses. arXiv preprint number 1009.1073, 2012.
- [108] Piotr Paweł Deuar. Private communication.