

Fuzzy Chaos: Quantum Particles and Classical Distributions

Robin Blume-Kohout

October 15, 1999

Abstract

I present an overview of the problems associated with quantum chaos, the need for a theory of semiclassical chaotic dynamics, and the importance of phase space for quantum systems. A selection of concepts are presented from each of: classical dynamics, quantum chaos, the mechanics of classical distributions, the quantum measurement problem, and open quantum dynamics. Finally, a candidate theory, continuous measurement, which permits semiclassical chaotic trajectories (and, in fact, classical behavior in the semiclassical regime) is presented. Appendices discuss classical phase-space distributions and the Wigner distribution for quantum systems.

1 Introduction

1.1 What's this paper about?

This paper is concerned with the semiclassical mechanics of single-particle systems, with particular emphasis on chaotic systems. The literature on chaotic systems is very extensive, so to avoid the necessity of either (a) quoting results without any of the background needed to understand them or (b) attempting to summarize an extraordinarily complex subject in a few pages, I will avoid rigorous mathematical discussions of semiclassical chaos. Instead, I will attempt to present the major problems in the field, some of the ideas that have been proposed to deal with them, and a few ideas that are extremely speculative and thought-provoking. This includes some preliminary theoretical and computational results from my own research into the subject.

1.2 What am I going to discuss

This paper is intended to introduce you, the reader, to some rather far-flung concepts in mechanics (classical and quantum, and statistical), which are tied together by *chaos*. To this purpose, we will begin by discussing chaotic behavior, or *nonlinear dynamics* for a bit. We will then consider the problem of *quantum chaos* briefly, though by no means exhaustively! Since one recurring topic is the connection between the mechanics of a statistical ensemble of classical particles and the mechanics of a quantum particle, we also discuss the mechanics of *distributions* of particles – with particular attention to the differences between particles that interact with each other and particles which don't interact with each other.

In particular, we'll examine semiclassical mechanics using wavepackets, and see how the Wigner distribution gives us a framework for connecting Hilbert spaces and phase spaces. Finally, we'll move back to quantum mechanics and discuss the problems with quantum measurement theory as motivations for semiclassical theories. This concludes the review section of the paper, after which I explain one candidate theory, *continuous measurement* dynamics.

2 Nonlinear dynamics

In the preface to his book, **Chaos in Classical and Quantum Mechanics**, Martin Gutzwiller comments that “Elementary mechanics, both classical and quantum, has become a growth industry in the last decade.” The “discovery” of (and subsequent explosion of interest in) chaotic systems in the 1960's bears a large portion of the responsibility for this growth. In the aftermath of relativity and the quantum revolution, classical mechanics had been considered a backwater, a field which was theoretically nearly complete. The discovery of chaotic behavior in very simple dynamical systems brought classical mechanics back into the limelight as physicists and mathematicians sought to quantify behavior that few scientists (with the notable exception of Poincare) had explored before.

2.1 Classical Dynamics: The High Points

Since I intend to do some rather unorthodox things with the concepts of mechanics, let's review those concepts quickly.

Consider a mechanical system \mathcal{S} (that is, a system governed by classical mechanics). The simplest system of any interest that we can imagine is

a point particle moving in a space of N dimensions, where N is usually 1, 2, or 3 in our universe. The state of this system \mathcal{S} at any time t can be completely described by its *position* (a set of N real numbers) and its *velocity* (another set of N real numbers that tell us how the position is changing). Instead of the velocity, we find it convenient to keep track of the *momentum*; thus the entire state of the system is summed up in two N -dimensional vectors, \vec{q} (its position) and \vec{p} (its momentum). For ease of notation (though perhaps not comprehension) most authors choose to lump these together into a single $2N$ -dimensional vector \vec{z} , so that the structure of \vec{z} is $\vec{z} = (q_1, \dots, q_N, p_1, \dots, p_N)$. For the particular case of a point mass moving in 3 dimensions, the explicit form of the state is $\vec{z} = (x, y, z, m\dot{x}, m\dot{y}, m\dot{z})$.¹ This notation can be extended to any mechanical system with a finite number of degrees of freedom, by letting N equal the number of degrees of freedom instead of the number of spatial dimensions. Any statements that we make about this formalism, then, will apply to a system with arbitrarily large or small N (unless an exception is specifically noted), but for the purposes of this paper, we will almost always consider N equal to 2 or 3.

Now, the fact that we can write down $2N$ real numbers and describe the “state of \mathcal{S} ” is a meaningless definition unless we flesh it out with some facts. Let us do so. First of all, the N q ’s and N p ’s define an abstract vector space that we call *phase space*. Since the state of \mathcal{S} can, at any time, be described by a $2N$ -dimensional vector in this phase space, we can imagine making many measurements of \mathcal{S} ’s state \vec{z} and constructing a set $\mathcal{T} = \{\vec{z}_i\}$ of all the states we measure. \mathcal{T} (which has a total ordering determined by the time of the measurement) is a *trajectory* in phase space, and if we take each $\vec{z} \in \mathcal{T}$ and get rid of the momentum part, then we’ve got the familiar coordinate-space trajectory that every student of freshman mechanics is familiar with. Everything I’ve said so far is familiar territory for any physicist, but I’m stating it explicitly because we’re going to need a firm foundation in a little while.

Now, in classical mechanics, the time-dependence of \mathcal{S} is contained in the Hamiltonian, H . That is, if we know the initial state of the system $\vec{z}(t = 0)$, then the trajectory \mathcal{T} that follows from that initial state can be computed by solving a system of first order differential equations:

$$\frac{d\vec{z}}{dt} = J \cdot \frac{\partial H}{\partial \vec{z}} \quad (1)$$

where J is a $2N \times 2N$ matrix defined in terms of the $N \times N$ null and identity

¹Well, as long as the system is subject to no forces which cannot be derived from a scalar potential. If a vector potential exists, $\vec{p} \neq m\vec{v}$.

matrices:

$$J = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix} \quad (2)$$

It is possible for the *form* of the Hamiltonian to depend upon the time, either periodically or stochastically. In the former case, some of the things I'm going to say are invalid; in the latter case, almost all of them are. For the duration of this paper, then, we will only consider time-invariant Hamiltonians. With this, we've set the formal foundation to talk about chaos.

The preceding formulation seems to make everything quite simple – after all, differential equations are not always easy to solve, but at least we have a formal solution to the whole problem of \mathcal{S} 's behavior. Unfortunately, things begin to get sticky. The Hamiltonian is a function of the N coordinates and the N momenta – for instance, for a *simple harmonic oscillator*, $H = \frac{1}{2}(\vec{p} \cdot \vec{p}/m + k\vec{q} \cdot \vec{q})$. While the formalism tells us that we can, in principle, derive an entire trajectory from a single initial condition, the differential equations are in practice often insoluble. The SHO happens to be easily soluble – but even such a simple modification as the addition of a q^4 term makes the equations insoluble. For many problems, we can obtain an approximate solution by *linearizing* the equations – that is, approximating high powers of \vec{p} and \vec{q} with linear and (sometimes) quadratic terms. This is, however, only an approximation, and in that fact lie the roots of chaos.

2.2 Emergence of Chaotic Behavior

We can divide dynamical systems into two categories: those with bounded phase spaces and those with unbounded phase spaces. The latter include particles in spaces free of force (which simply sail at constant velocity off to infinity) and particles in constant force fields (which, if we ignore relativity, accelerate off to infinity at constant acceleration). In general, these systems are either uninteresting or hard to deal with; although some very interesting systems (the 3-body problem, for instance) are unbounded, we can almost always define a bounded region of phase space which contains all the interesting motion. Once the system leaves this region, its subsequent behavior becomes of little interest. Thus, let's focus on systems with bounded phase spaces. If the phase space has limited volume, then the coordinate (or *configuration*) space is similarly limited in volume, and unless the system has a way of losing energy and coming to rest, it must oscillate or jiggle around in this limited volume. We'll also confine ourselves, for now, to *Hamiltonian*, or non-dissipative, systems – that is, systems which have no means of losing

energy. Thus, we are restricted to systems which oscillate in one way or another for our study of chaos. This is not an arbitrary choice on my part, by the way; oscillation is a universal characteristic of non-dissipative chaotic systems.

As shown previously, mechanics provides us with a way of predicting the state of a system at some time t , given only its [time-independent] Hamiltonian and its state at time $t = 0$. However, this presumes that we “know the state at time $t = 0$ ”. To “know” the state in this case means specifically to know it *exactly* – an impossible proposition from everything we know of physics. No meterstick can measure the position of a particle infinitely well; no camera can record its velocity infinitely well. Thus the entire premise that our theory of mechanics is founded on (so far) is extremely shaky.

The obvious question is: How much of a difference does a small uncertainty in the initial state $\vec{z}(0)$ make? If we know the initial state to within 1%, do we know the state 10 minutes later to within 1%?

To answer this, we resort to perturbation theory. If we know an initial state $\vec{z}(0)$ to within $\pm\delta\vec{z}(0)$, and we need to know how much uncertainty this provides in the final state $\vec{z}(t)$, then it is reasonable to look for some equation of motion governing $\delta\vec{z}$. Such an equation exists. There exists a definite trajectory corresponding to the initial state $\vec{z}(0)$, which we can call $\vec{z}(t)$. For arbitrarily small initial $\delta\vec{z}(0)$, $\delta\vec{z}$ evolves according to:

$$\frac{d}{dt}\delta\vec{z} = J \cdot H''(t) \cdot \delta\vec{z} \quad (3)$$

where H'' is the $2N \times 2N$ *Hessian* matrix of the Hamiltonian H :

$$H''_{\alpha\beta} = \frac{\partial^2 H}{\partial z_\alpha \partial z_\beta}(\vec{z}(t)) \quad (4)$$

What does this tell us about particular systems? When applied to many systems, we find that the uncertainty in our knowledge of \mathcal{S} 's state either remains constant as time evolves or increases linearly in time. Incidentally (because of Liouville's theorem) it never decreases unless we know at least one of the initial state variables to infinite precision – which is impossible. The systems which display this sort of *linear* behavior are called *linear systems*, because their equations of motion typically contain only terms that are linear in the state variables.

When we apply this analysis to other systems, however, we find that the uncertainty $\delta\vec{z}(t)$ grows **exponentially** in time! That is, if we know the initial state to D significant figures, then after an amount of time that is

only logarithmic in D , our analysis of \mathcal{S} 's behavior is essentially meaningless, because the uncertainty has grown to overwhelm the actual result. That is, $\delta\vec{z}(t_f)$ and $\vec{z}(t_f)$ are on the same order of magnitude for some t_f which is proportional to $\log(D)$. Contrary to a common misconception, this does *not* mean that the uncertainty in \vec{z} grows as e^t for all time – because the phase space of the system is almost always bounded and/or periodic, there is a hard limit on $\delta\vec{z}$. The exponential growth is only valid when the uncertainty is significantly smaller than the state itself.

This phenomena is known as “sensitive dependence on initial conditions,” which means exactly what it sounds like. The final state of the system depends so sensitively on the initial state of the system that long-term prediction is essentially impossible – no matter how well we define the initial state. This *chaotic* behavior is very different from *stochastic* behavior; the dynamics of a stochastic system are governed by random dynamical processes, such as Brownian motion, which imply a time-dependent Hamiltonian. A chaotic system can have a simple, time-independent Hamiltonian, yet display complex motion that cannot be predicted beyond a short time in the future.

This is essentially where my exposition of chaotic dynamics must stop. The subject is entirely too rich to cover in any more depth; however, I would like to present three brief topics.

- First, let us consider (briefly) some basic theorems about chaotic dynamics. Recall that phase space is a $2N$ -dimensional vector space. We can define a *phase space volume* \mathcal{V} by means of $2N$ linearly independent vectors in phase space, and we can then consider the set of all states $S = \vec{z}$ that lie within the volume \mathcal{V} . *Liouville's Theorem* proves that if we use the time evolution of the states of our system \mathcal{S} to create a map from each point in phase space at $t = 0$ to some point in phase space at $t = t'$, and look at what effect this has on the set S , then the phase space volume \mathcal{V}' occupied by S at time t' is equal to S 's original volume \mathcal{V} . A concise statement of Liouville's theorem is that the set of all phase space orbits evolves in time like an incompressible fluid (for Hamiltonian systems).

For there to be chaotic dynamics and sensitive dependence on initial conditions (SDIC) at all, though, there must be at least one dimension in which the set S stretches out exponentially in time (so that $|\delta\vec{z}|$ can increase exponentially). To preserve volume, if the volume stretches out along one dimension, it must contract along some other dimension. This behavior, the stretching and contracting of phase space volumes

over time, is defined by the *Lyapunov exponents* of the point in phase space around which S is centered. S (technically, the point around which it lies, but if \mathcal{V} is small it amounts to the same thing) has $2N$ Lyapunov exponents L_α , which characterize the deformation of S over time in the following simple, intuitive way. If S initially occupies a $2N$ -dimensional hyperrectangle with side lengths l_α in phase space, then as time goes on, each side length l_α is stretched or contracted proportional to $\exp[L_\alpha t]$. The volume of this hyperrectangle is thus just

$$\mathcal{V}' = \mathcal{V} \prod_{\alpha=1}^{2N} \exp[L_\alpha t] = \mathcal{V} \exp \left[t \sum_{\alpha=1}^{2N} L_\alpha \right] \quad (5)$$

In order that phase space volume be preserved (as Liouville's theorem requires), the sum of the Lyapunov exponents must be zero. Chaotic systems have trajectories in their phase spaces along which at least one of the $2N$ Lyapunov exponents is positive. A trajectory along which all the Lyapunov exponents are zero is a periodic (nonchaotic) orbit – regardless of whether the system in which it occurs displays chaos otherwise or not. Incidentally, even if all the Lyapunov exponents for a trajectory are zero, a volume of phase space which follows that trajectory will still (usually) be deformed. If it expands or contracts *linearly* along one dimension \hat{z}_α , the Lyapunov exponent L_α is still zero.

- Second, we might well wonder just how such “chaotic” behavior can occur. This is particularly relevant to the discussion in the later sections of this paper, so let's consider it. The dynamics of a mechanical system can be thought of as an *iterative* process, in that any trajectory can be broken down into many very small segments, each one of which is approximately linear. That chain of segments (in the limit as the segments get very small and very linear) is an iterative process; each successive point along the trajectory can be obtained from the previous point by a simple linear operation. If we let \hat{N} (for **N**ext) be the operator that takes a point \vec{z}_n to the next point \vec{z}_{n+1} , then we write $\vec{z}_{n+1} = \hat{N}\vec{z}_n$, and in general, $\vec{z}_{n+m} = \hat{N}^m \vec{z}_n$. Thus we have a prescription for evolving the system over any time period; simply apply \hat{N} the right number of times.

Now, let's consider some iterative systems and the corresponding operators. A simple iterative system in one dimension might be:

$$x_{n+1} = x_n + 1 \quad (6)$$

This isn't particularly exciting; it just leads to x increasing linearly in time. If we look at the behavior of two different x_0 's, separated by a small amount, then their difference remains constant in time. Clearly, if the *effect* of the operator is independent of x_n (in this case, \hat{N} is the “add-one-to-the-argument” operator), then the behavior is predictable – *linear*. A similar case is the system:

$$x_{n+1} = e^{i\phi}x_n \tag{7}$$

This is a bounded system; x simply rotates in the complex plane. Again, the behavior of the “rotate-by- ϕ ” operator is independent of the value of x_n . However, if we arrange that *the effect of the iteration operator is highly dependent on the value of x_n* , then we find a very different result, as in:

$$x_{n+1} = a(1 - x_n)x_n \quad (0 < x_0 < 1), \quad a > 1 \tag{8}$$

Now the operator is nonlinear in x_n , and we find (by experiment) that two very nearly identical values of x_0 lead to exponentially diverging trajectories if a is sufficiently large.

What does all this imply about mechanical systems? Simply that chaotic behavior is the result of having equations of motion *whose effect, as iteration operators, depends on the state on which they act*. This is what is meant by non-linear behavior; to get from point A to point B you have to apply an iteration operator lots of times, and the nonlinear terms in that operator get magnified when you apply it over and over. So any dynamical scheme in which the dynamics is dependent on the state will display exponential behavior of one sort or another; in stable oscillators the exponential behavior is a result of negative feedback, leading to a complex exponential (oscillation); in chaotic systems the exponential behavior is a result of positive feedback, which leads to a real-valued exponential (sensitive dependence on initial conditions). We'll want to consider this concept again later.

- Finally, what does an information-theoretic analysis of chaos tell us? This point is of no particular use for the development of the theory, but it does provide a motivation for exploring these questions. Information theory is concerned with the nature of information, its transfer, and its production and destruction. Recently, attempts have been made to formulate physics entirely in terms of information theory.

The information-theoretic significance of chaos is as follows. If we know the initial conditions of a system to within, say, 4 significant figures, then this is a certain amount of information. Four digits corresponds to one part in 10,000, or approximately 13 bits of information. If the system dynamics are linear, then the uncertainty in our knowledge of the system grows at most linearly in time. At $t = 0$, we know z to within 0.01%; at $t = 1$ we know it to within 0.02%; at $t = 100$ we know it to within 1%. If we have a measuring apparatus that can make measurements with an accuracy of 1 part in 10^4 , then by making another measurement at $t = 100$, we can reduce our uncertainty at $t = 100$ by a factor of 100. This means we can backtrack (run the evolution equations backwards) to $t = 0$ and know the initial state to within a factor of 10^6 . In short, the time evolution of the system has produced approximately 7 more bits of information! This concept of information production is essential in information theory; we never have unlimited information about the universe, and any process that increases the information which we have about the universe is significant. However, we can see that our linear system only produces information logarithmically – that is, every time the elapsed t increases by a factor of 2, we get another bit of information. For a system whose uncertainty is constant in time, the rate of information production is nil. In either case, the amount of information that the system can produce over, say, the age of the universe is strictly limited.

A chaotic system, however, has the property that the uncertainty in z increases exponentially. A moment's thought will show that this leads to a linear production of information – if δz doubles every timestep, then another bit of information is produced every timestep. Of course, the question of whether we *care* about the information that we're getting is still open, as is the question of how efficiently we can collect the information. Theoretically, however, the system has an infinite information capacity (read Stanislaw Lem's Cyberiad for a fascinating science-fiction story on this phenomenon). This would only be a curiosity were it not for the fact that *quantum chaos* is so elusive. The non-existence of chaotic information-producing systems in quantum mechanics poses some interesting problems for quantum information theory (and quantum computing), but any more discussion of this phenomena belongs in another paper.

3 Quantum Chaos

The theory of quantum mechanics was developed for the particular purpose of explaining the behavior of very small or low-energy objects – atoms, photons, subatomic particles. Like the theory of special relativity, the quantum theory’s very existence demonstrates that classical mechanics is in some fashion “incorrect”; classical mechanics yields the wrong results when applied to certain systems. Relativity, however, is something of a model theory in that relativistic calculations of phenomena in the “classical regime” yield results that are identical with those derived from classical mechanics. The only exceptions to this rule are those where the relativistic result is demonstrably correct – for instance, magnetism, a phenomenon which lies apparently within the “classical regime,” yet is actually due to relativistic effects. Thus, classical mechanics and Galilean relativity may be said to fit seamlessly into the framework of relativistic mechanics.

Quantum mechanics does not fit seamlessly into the framework, however. There are many demonstrations of this, some of which we shall explore later, but for now let us consider the particular case of chaotic dynamics in quantum systems.

There are several signatures of chaos in dynamical systems (some of which have been discussed above); we should pick one of them as a bellwether to identify chaos in quantum systems. The most fundamental (this is open to debate) seems to be *sensitive dependence on initial conditions* – that is, the existence of at least one positive Lyapunov exponent along a trajectory. Testing this seems straightforward; we can prepare two initial states which are in some sense “close,” then let them evolve and see how they diverge. Let’s examine this process more thoroughly.

The time-dependent version of the Schroedinger equation is

$$i\hbar \frac{\partial \psi}{\partial t} = \hat{H} \psi \tag{9}$$

Let’s start with two states $|\psi(0)\rangle$ and $|\xi(0)\rangle$. We need a way to define “closeness,” and the most natural way to ensure that $|\psi(0)\rangle$ and $|\xi(0)\rangle$ are “close” is to specify $\langle \psi(0) | \xi(0) \rangle = \alpha$, where α is nearly 1. $\langle \psi(0) | \xi(0) \rangle$ is a measure of closeness in Hilbert space, so this seems an appropriate way to measure closeness.

We can write $|\psi(0)\rangle$ and $|\xi(0)\rangle$ as expansions in an energy eigenbasis:

$$\begin{aligned} |\psi(0)\rangle &= a_1 |u_1\rangle + a_2 |u_2\rangle + a_3 |u_3\rangle \dots \\ |\xi(0)\rangle &= b_1 |u_1\rangle + b_2 |u_2\rangle + b_3 |u_3\rangle \dots \end{aligned} \tag{10}$$

Applying Eq. 9 to Eqs. 10, we get:

$$\begin{aligned} |\psi(t)\rangle &= a_1 e^{-i\frac{E_1}{\hbar}t} |u_1\rangle + a_2 e^{-i\frac{E_2}{\hbar}t} |u_2\rangle + \dots \\ |\xi(t)\rangle &= b_1 e^{-i\frac{E_1}{\hbar}t} |u_1\rangle + b_2 e^{-i\frac{E_2}{\hbar}t} |u_2\rangle + \dots \end{aligned} \quad (11)$$

But now, when we calculate the overlap of $|\psi(t)\rangle$ and $|\xi(t)\rangle$, we find that $\langle\psi(t)|\xi(t)\rangle = \alpha = \langle\psi(0)|\xi(0)\rangle$. The overlap has not changed! This is the fundamental problem of quantum chaos. Quantum states exist naturally in a Hilbert space, but because the time evolution operator is unitary, it preserves Hilbert space inner products. Thus two states that begin life close together necessarily end life close together, and divergence of any kind (much less exponential divergence) is impossible.

There are, not surprisingly, many responses to this phenomena. One of the most generally valid is to note that Hilbert space and phase space are most definitely not the same thing. They're both vector spaces, but other than that there are really very few similarities. Thus, to equate separation in Hilbert space (in which, because vectors are normalized, separations are equivalent to angles) with separation in phase space (in which vectors are not normalized), is somewhat disingenuous. One problem we would like to be able to address, therefore, is how to put classical and quantum *states* (never mind their dynamics for now) on an equal footing.

Additionally, we ought to consider the fact that when we consider two vectors separated by an angle θ_α in each dimension Θ_α of any normed vector space, their inner product is equal to the product of the cosines of the θ_α 's. Thus, as the number of dimensions in the space goes to infinity, the inner product of any two states separated by a small amount in each dimension approaches zero, and we can assume that any small random perturbation on a state in an infinite-dimensional Hilbert space will yield a state whose overlap with the initial state is virtually zero. The problem with this argument is that while it applies to quantum states of macroscopic objects, which have almost infinitely many independent degrees of freedom (for instance, the degrees of freedom corresponding to non-adjacent atoms in a molecule are almost independent), it doesn't apply to the infinite-dimensional Hilbert spaces of microscopic systems.

Consider, for example, a single free neutron². In the position basis, there are infinitely many degrees of freedom; thus any small perturbation on the neutron state should yield another state which is essentially orthogonal to the first. However, each eigenstate of position $|x_1\rangle$ is coupled to every other

²We ignore the spin dimension

$|x_2\rangle$ by the momentum and energy operators. Thus any small perturbation which takes $|\psi_1\rangle$ to $|\psi_2\rangle$ and leaves $\langle\psi_1|x|\psi_1\rangle \simeq \langle\psi_2|x|\psi_2\rangle$ while ensuring $\langle\psi_1|\psi_2\rangle \simeq 0$ will almost undoubtedly change the momentum of the state drastically, so $|\langle\psi_1|p|\psi_1\rangle| \ll |\langle\psi_2|p|\psi_2\rangle|$. The 5-cent explanation for this is that \hat{x} and \hat{p} span the operator space of the free neutron completely, so any two states which are “nearby” in the space of one operator yet have negligible overlap must be very far apart in the space of the other. This ties in closely with the Wigner function, which we’ll discuss later.

There seem to me to be (and this is open to debate) three main problems in the theory of quantum chaotic systems in phase space. Before I enumerate them, please note that we are specifically considering quantum chaos *in phase space*. Much of the work on quantum chaos has addressed the eigenfunctions of quantized chaotic systems (most notably Heller’s discovery of scars), their spectra (level spacing distributions and random matrix theory, among other areas), and the dynamics of systems with low Hilbert space dimensionality (the hydrogen atom in a magnetic field, spin-flip systems, and a host of others). Physical chaos, however, is really a feature of classical mechanics more than anything else³, and classical mechanics is naturally dealt with in phase space. Thus, while mathematical rigor is more difficult to achieve in a treatment of quantum chaos based on phase space, any progress is deeply satisfying, since it bears directly on the correspondence between quantum and classical mechanics.

1. First, how do we pose quantum mechanics in phase space instead of Hilbert space? This is a fascinating area, but because the connection between classical and quantum mechanics has long been of interest to physicists and mathematicians, it is a thoroughly mined area as well. Feynman’s path-integral picture is generally perceived as being of the most interest; the concept of a path is classical at heart, and it’s not hard to express path integrals as integrals over paths in phase space. In addition, that paradigm is often useful for doing calculations and proofs in semiclassical mechanics. The other particularly useful connection between quantum mechanics and phase space is the idea of distributions in mock phase space – essentially, a quantum mechanical phase space distribution. The most general of these is probably the Wigner function.

2. In a more general vein, how can we obtain nonlinear dynamics from a quantum theory in which time evolution is inherently unitary and

³Mathematical chaos is a different beast.

linear? The Schroedinger equation is both of these. The linearity of time evolution seems to preclude nonlinear divergences in any system with relatively few macroscopic degrees of freedom, while unitarity precludes the information production that (as we saw in our discussion of chaos) classically chaotic systems demonstrate. This is a much more recent field of inquiry, and the basic solution to both of these problems is to attempt to treat in some coherent way the inevitable interaction of the “system” with the “environment.” These terms are in quotes for a reason: the distinction is rather arbitrary. In practice, separation of system and environment is difficult, because while we’d like the “system” to be simple and isolated, we can neither impose a potential on it nor make any measurement of it without introducing something that ought to be considered “environment.” This leads to the theory of open quantum systems, which includes among other things alternative formulations of time evolution. This is the subject I am particularly eager to address.

3. Finally, what are we going to do about measurements? In classical mechanics, we recall that it is possible to determine the time evolution of a system’s state entirely by means of a Hamiltonian, an initial state, and a set of differential equations. This theoretical trajectory can then be compared with experimental measurements which (in theory) are made without disturbing the system at all. That is, the system should evolve according to our mathematical model regardless of whether we make measurements or not. In quantum theory, of course, this breaks down entirely; we can calculate analytically the time evolution of the *quantum state* for all time, but the quantum state does not have a one-to-one mapping with measurements! Instead, in order to actually get any data from our system, we must disturb it by making a measurement. Thus, while a formulation of quantum dynamics which results in wavefunctions that follow classical paths is satisfying, we must consider the measurement process in order to have anything resembling a complete theory of quantum dynamics (chaotic or not).

We now have something of a plan for exploring avenues towards semiclassical chaos⁴ in phase space. We should first attempt to form an idea of how

⁴I attempt to maintain the distinction between “quantum chaos” and semiclassical chaos” in the following (somewhat artificial) manner. *Quantum* chaos, for the purposes of this paper, is considered to be any manifestation of chaotic, aperiodic, or nonlinear behavior whose signature is uniquely quantum mechanical, and which does not rely on

quantum mechanics relates to phase space, then try to find a formulation of quantum dynamics that satisfies four conditions.

1. It provides, in theory, for classically sensible trajectories which can diverge exponentially.
2. It includes some element of nonlinearity and/or nonunitarity, so that information is produced.
3. It takes into account (and accounts *for*) the measurement process by which we get information about the trajectory.
4. It agrees in the appropriate limits with unitary quantum mechanics and classical mechanics.

Needless to say, this is an agenda which (a) has not yet been fully realized and (b) is certainly not going to be completed in the course of this paper! However, we can examine some efforts to realize it.

4 Semiclassical Mechanics and Wavepackets

4.1 What is a Wavepacket?

The use of wavepackets to form a theory of semiclassical mechanics is very nearly as old as quantum mechanics. A wavepacket is simply a quantum state which is not, in general, an eigenstate of any operator (although it may be coincidentally an eigenstate of energy), but is in some fashion localized in position and momentum space. The most common shape for a wavepacket is a Gaussian, but we can imagine a wavepacket of any shape. Certain shapes are not “wavepackets” per se; for instance, a square wavefunction is theoretically realizable, but cannot be localized in momentum space; therefore it spreads infinitely fast in position, and is not a useful wavepacket. The point of wavepackets is that they permit realization of certain classical properties over finite times; under unitary (Schroedinger) dynamics, all wavepackets spread out in q -space over time, but if they are localized in p -space, then they can behave semiclassically until they spread beyond some limit.

any assumptions about \hbar , space/time scales, or approximated potentials. *Semiclassical* chaos is taken to refer to any theory which recovers from quantum dynamics a classically recognizable dynamical signature of chaotic behavior, particularly involving well-defined trajectories and phase space.

4.2 Properties of Gaussian wavepackets

As mentioned in the previous section, a Gaussian wavefunction realizes the minimum joint uncertainty for a quantum state.⁵ In addition, each Gaussian state is the eigenfunction of some harmonic (quadratic in q) Hamiltonian. If a Gaussian wavepacket is placed at the origin in a quadratic potential of which it is the eigenfunction, it will remain stationary; if, however, it is displaced from the origin, then the state maintains its shape while oscillating back and forth in the potential. For obvious reasons, such states are referred to as *coherent states*. Furthermore, a Gaussian wavepacket in a potential of which it does *not* happen to be an eigenstate maintains its overall shape, but oscillates in size as well as position. In general, a Gaussian wavepacket in any constant, linear, or quadratic potential remains a Gaussian in shape. Since a Gaussian is shape-invariant under the Fourier transform, we can describe the state of a Gaussian at any point in time with only $4N$ real numbers ($\langle \vec{q} \rangle$, $\langle \vec{p} \rangle$, $\Delta \vec{q}$, $\Delta \vec{p}$).

Like any wavefunction, all the properties except $\langle \vec{q} \rangle$ and $\langle \vec{p} \rangle$ of Gaussian wavepackets are invariant under Galilean transformations. This means that we can change reference frames either in q -space (through spatial translation) or in p -space (by multiplying the wavepacket by a phase factor $\exp[i\vec{k} \cdot \vec{q}]$) without changing any of the properties of the wavepacket. A more useful formalism for these transformations (but one which we have no space to discuss here) is the set of *Heisenberg operators*, which move wave functions around in position and momentum space. Unlike Galilean transformations, the Heisenberg operators display the non-commutativity geometry of quantum phase space – that is, if the same wavepacket is translated along two different paths to the same point, it will pick up different phases along the two paths depending on their length. This leads to some very interesting mathematics; the interested reader is referred to the reference at the beginning of the next section.

4.3 Formulation of Semiclassical Mechanics using Gaussian Wavepackets

The interested reader is immediately referred to R.G. Littlejohn's excellent review article, "The Semi-Classical Evolution of Wavepackets" in Physics Reports v. 128 n. 4 (May, 1986). What is presented here is primarily distilled from that extensive survey of wavepacket mechanics.

⁵The corresponding Wigner function therefore occupies the minimum allowable volume for a quantum state in phase space

Wavepackets are a particular attractive gateway to quantum/classical connections because they can be understood in so many ways. As long as the local potential experienced by a wavepacket $\psi(\vec{q})$ is quadratic or lower-order in \vec{q} , ψ remains a wavepacket. This can be shown through the Schroedinger equation, through the path-integral formulation of quantum mechanics, and from the WKB theory. In addition, Ehrenfest's theorem⁶ applies rigorously to wavepackets moving in quadratic potentials. By this I mean that not only are the general Ehrenfest relations:

$$\frac{d\langle\hat{q}\rangle}{dt} = \frac{\langle\hat{p}\rangle}{m} \quad (12)$$

$$\frac{d\langle\hat{p}\rangle}{dt} = -\left\langle\frac{\partial V}{\partial\vec{x}}(\hat{q})\right\rangle \quad (13)$$

satisfied, but the stronger relations:

$$\frac{d\langle\hat{q}\rangle}{dt} = \frac{\langle\hat{p}\rangle}{m} \quad (14)$$

$$\frac{d\langle\hat{p}\rangle}{dt} = -\frac{\partial V}{\partial\vec{x}}(\langle\hat{q}\rangle) \quad (15)$$

are true. The difference is that equations 12-13 state that the expectation values of the quantum observables evolve according to the canonical equations of motion derived from a Hamiltonian whose potential energy is averaged over the wavepacket, while equations 14-15 indicate that the potential in the effective classical Hamiltonian is evaluated at the center of the wavepacket $\langle\vec{q}\rangle$. Thus wavepackets in quadratic potentials follow the classical motion exactly.

In general, if the potential experienced by a wavepacket is not uniformly quadratic, then we can make a Taylor expansion of the potential in the region of the wavepacket to see if it is approximately quadratic over the [fairly well-defined] radius of the wavepacket. Essentially, because the wavepacket does not feel the potential outside of its own radius, we can assume that it feels an effective quadratic potential as long as the scale on which cubic and higher-order terms in $V(\vec{q})$ become significant is much larger than the radius of the wavepacket. This *local quadratic approximation* is not merely a quantum mechanical technique; it appears in classical perturbation theory, and we will shortly use it classically. In passing, we note that in certain physical systems, quadratic approximations in momentum are also necessary; for mechanical systems, however, the Hamiltonian is always quadratic in \vec{p} .

⁶Ehrenfest's theorem states that the expectation values of quantum observables evolve classically – that is, in parallel with the values of classical observables

Gaussian wavepackets have one final property of particular interest. We can construct a free-particle wavepacket which has minimal $\Delta q \Delta p$ at $t = 0$, but as time passes, the wavepacket spreads out in \vec{q} without contracting in \vec{p} , so that it no longer has minimal uncertainty. It is also possible to construct a wavepacket which does not have minimal uncertainty. However, it can be shown any wavepacket $\psi(\vec{q}, t)$ has minimal uncertainty at some time t_0 which may be either in the past or the future. Thus there is not only a minimum on the uncertainty in \vec{q} and \vec{p} for a quantum state, but also a sort of maximum; every wavepacket is identical to every other wavepacket under two Galilean transformations in \vec{p} and \vec{q} and a time evolution!

4.4 Spreading of Gaussian Wavepackets: Nearby Orbits and “Quantum Diffusion”

The dynamics of a free-particle Gaussian wavepacket in one dimension can be solved exactly, and the solution generalized to three or more dimensions. This analysis also applies to motion in a scalar potential, but because the mathematics is less simple and attractive, we examine only the free-particle case.

For a free particle, the Hamiltonian is a function only of \vec{p} , and so momentum is a constant of the motion. The coordinate \vec{q} is not, however, and so the wavepacket expands or spreads as time passes. In the limit of long time, this spreading is linear in time; we can show by explicit computation with the Schroedinger equation that if the initial state of the wavepacket is given by:

$$\psi_0(x) = \frac{1}{(2\pi L^2)^{1/4}} \exp\left(-\frac{x^2}{4L^2} + \frac{ip_0x}{\hbar}\right) \quad (16)$$

where $\Delta x_0 = L$ and $\Delta p_0 = \hbar/2L$, then the state of the packet at time t is:

$$\begin{aligned} \psi(x, t) = & \frac{1}{(2\pi L^2)^{1/4}} \frac{1}{\sqrt{1 + \frac{i\hbar t}{2mL^2}}} \cdot \\ & \exp\left[-\frac{\left(x - \frac{p_0 t}{m}\right)^2}{4L^2 \left(1 + \frac{i\hbar t}{2mL^2}\right)} + \frac{ip_0x}{\hbar} - \frac{ip_0t^2}{2m\hbar}\right] \end{aligned} \quad (17)$$

From this it follows that the expectation values and dispersions of the state variables at time t are:

$$\langle \hat{x} \rangle = \frac{p_0 t}{m} \quad (18)$$

$$\langle \hat{p} \rangle = p_0 \quad (19)$$

$$\Delta x(t) = L \sqrt{1 + \frac{\hbar^2 t^2}{2m^2 L^4}} \quad (20)$$

$$\Delta p(t) = \Delta p_0 \quad (21)$$

The Ehrenfest relations are satisfied in this simple case. In addition, we see that the spreading of the wavepacket is quadratic for short times and linear for long times. The appearance of \hbar in the expression for the wavepacket width is often interpreted to mean “this is a strictly quantum effect, and disappears for classical systems”, but this interpretation is not entirely correct. This spreading is commonly referred to as “quantum diffusion,” another misnomer.

To see this, we consider a Liouville ensemble $f(x, p, t = 0)$ given by the Gaussian distribution:

$$f_0(x, p) = \frac{1}{2\pi LK} \exp \left[-\frac{x^2}{2L^2} - \frac{(p - p_0)^2}{2K^2} \right] \quad (22)$$

We can calculate the distribution at a later time t by a simple substitution:

$$\begin{aligned} f(x, p, t) &= f\left(x - \frac{pt}{m}, p\right) \\ &= \frac{1}{2\pi LK} \exp \left[-\frac{(q - pt/m)^2}{2L^2} - \frac{(p - p_0)^2}{2K^2} \right] \end{aligned} \quad (23)$$

The expectation values evolve as expected, and we find that the variance, or width of the distribution at time t is:

$$\Delta q(t) = L \sqrt{1 + \frac{K^2 t^2}{m^2 L^2}} \quad (24)$$

First of all, the functional form of equation 24 is identical to that of the quantum packet’s width, equation 20. In addition, we can interpret the factor of \hbar in the previous expression in the following way. If we hold everything in the first expression fixed and let $\hbar \rightarrow 0$, then the spreading does go to zero. However, that corresponds to localizing a quantum state to a volume of phase space less than $\frac{1}{2}\hbar$, which is impossible. If, on the other hand, we choose our classical distribution so that it initially occupies a volume of phase space equal to $\frac{1}{2}\hbar$, then we have $LK = \frac{1}{2}\hbar$, and substituting this into equation 24, we find that the spreading rate is identical to that given by equation 20.

The incorrectness of the term “quantum diffusion” is now apparent; first, diffusion is a stochastic process, but we have not invoked stochasticism or the Fokker-Plank equation here. The diffusion equation is second-order in space and leads to a spreading rate that goes as \sqrt{t} ; we have both quantum and classical models which are first-order in space and spread linearly in time for large t . The obvious conclusion is that there is a classical interpretation not only to the evolution of wavepacket expectation values, but also to wavepacket spreading. Our success in this case indicates that there may be other ways to comprehend the spreading of wavepackets in classical terms. The previously mentioned review article by Littlejohn pursues this in much greater mathematical detail. For now, however, having found a simple model which connects quantum and classical distribution dynamics, we move on to seek manifestations of semiclassical chaos.

5 Quantum Mechanical Issues

I previously suggested four goals for a theory of semiclassical chaos. We can see how the semiclassical theory of wavepackets fulfills at least part of the first goal (classically sensible trajectories, a classical analogy to wavefunction spreading, and appropriate limiting behavior), but in order to obtain exponentially diverging trajectories, information production, or an account of the measurement process, we must alter our picture of quantum dynamics. First, let’s take a whirlwind tour of what might be described as “Alternative Quantum Mechanics.”

5.1 The Quantum Measurement Problem

As I mentioned initially in discussing quantum chaos, quantum theory fails to agree in the appropriate limits with classical mechanics. In particular, there are aspects of quantum dynamics that cannot be smoothly scaled in any parameter (as relativity can be scaled in γ) to match classical behavior. Foremost among these is the process of quantum measurement – the Collapse Postulate.

The familiar picture wherein a measurement of an observable \hat{A} produces an eigenvalue a of \hat{A} and leaves the measured system in an eigenstate $|a\rangle$ of \hat{A} is derived from nothing; it exists as an independent postulate in quantum theory. Perhaps no other single physical postulate has provoked so much argument and frustration. The collapse postulate *works*, yet it provides so many conceptual problems that it’s difficult to know where to

begin. First of all, the measurement of a continuous state variable (position, for instance) to infinite precision requires an infinitely long time if the measured quantity commutes with energy and infinitely high energies if it does not. In practice, quantum systems exist on energy and distance scales such that these objections are purely theoretical; nonetheless, eigenstates of position are necessarily unphysical. Second, the measurement of one of a pair of noncommuting quantities should introduce infinite uncertainty in the conjugate quantity, yet this can also yield nonphysical results. Finally, it is unsatisfying to deal with a measurement theory which permits no description of uncertainty in measurements.

Many of these objections have been addressed in one way or another in practice over the past half-century. Quantum optics, in particular, has encouraged development of the idea of *quantum non-demolition measurements*. While no single theory has arisen to replace the collapse postulate, there are increasingly many alternatives to the collapse theory of measurement. Ideally, we would like to replace the collapse postulate with something which fits conveniently into the framework of quantum dynamics, instead of nailing it on *a priori*.

5.2 Open Quantum Systems and Weak Measurements

5.2.1 The Environment Problem and Open Quantum Systems

The other major problem in quantum mechanics that has led to the extension of quantum theory is the problem of system-environment interaction. This is a rather old field; the density matrix, for instance, was devised and explored in the early years of the quantum theory. The problem is rather simply stated; we would like to be able to deal with quantum systems which have a finite (and small) number N of degrees of freedom. However, any such system (for instance, a free particle) keeps on interacting with other systems. At best, we can make the contribution of these interactions to the overall Hamiltonian weak in comparison to the internal energies, but this is not always possible. A charged particle, for instance, is constantly interacting with the electromagnetic field, and thus with every other charged particle in the universe.

We are left with three choices: ignore the effects of the environment; let every system have infinitely many degrees of freedom; or develop a formalism for analyzing the effects of a general “environment” on a system of interest. The theory which developed from this dilemma is the theory of *open quantum systems*, and it finds particular application in quantum

computing, where the interactions of a quasi-isolated quantum system with *anything* are of paramount importance. There are a number of ways of formulating open system dynamics, perhaps the most general of which are the master equation approaches. The Lindblad equation, for instance, describes the time evolution of the density matrix representing a system in contact with an environment. Unlike the Schroedinger equation, however, the *form* of the Lindblad equation is not completely specified; just as the Schroedinger equation requires the Hamiltonian to be specified, the Lindblad equation requires a set of operators describing the effects of the environment. These are not well-defined, and may even be impossible to specify. Thus, other approaches are useful.

5.2.2 Weak and Fuzzy Measurements

A particularly simple model of open systems works as follows. We know that large portions of the universe which makes up the total environment of our system are very much in the classical regime. The interactions of quantum and classical systems usually result in measurements, since our measuring apparatuses are always classical. Thus, we can try to model the environment as some large machine which makes measurements on the system. At this point, however, the quantum measurement problem discussed above looms large, since if we obey the collapse postulate, every measurement of our system destroys all previously existing information in its state. A more general theory of measurement is necessary, and has fortunately been provided.

In general, we can make two sorts of measurements which are in some manner *weaker* than the sort of measurement which causes collapse. First, we can measure operators which may not correspond to a classical state variable, but which only yield a limited amount of information. In the case of the free neutron which we examined so long ago, we might make a measurement of position which only determined in which half of a room the neutron was to be found. The result would be a single bit of information, but half of the information in the wavefunction would be retained. This sort of measurement is called a **weak** measurement.

The problem with the measurement just described is that if the neutron is determined to be on, let us say, the left side of the room, then the probability for it to be found at that instant on the right side of the room is reduced immediately to zero. If the value of the wavefunction at the boundary between left and right halves of the room was significant at the time of measurement, then the post-measurement wavefunction has a discontinuity – which necessarily disturbs the value of momentum tremendously. The

basic problem here is that we are making a weak but very sharp measurement – there is no uncertainty about which side of the room the neutron occupies at the time of measurement. As a result, the induced uncertainty is very much greater than the ideal case.

The solution to this problem is quite physical and natural; we suppose that our measuring environment has only a finite energy at its disposal. The measurements made by such a machine will inevitably be somewhat unreliable – but we can remain assured that our measurement does not disturb the energy of the system by an amount greater than the energy available to the environment/measuring machine. Such an assumption produces a **fuzzy** measurement, and in conjunction with the restriction described previously, can produce a **weak, fuzzy measurement**. The environmental model which is most easily made rigorous (and physical) is an ensemble of harmonic oscillators of varying frequency, coupled weakly to the system. Descriptions of this sort of canonical ensemble are common in the literature of open quantum systems. Another interesting environmental model is a bath of low-energy photons which interact with the system through scattering. If the measurements of such an environment are carefully modeled to approach the minimal uncertainty of $\frac{1}{2}\hbar$, the effects on the dynamics of a system can be remarkable.

5.3 A Smidgen of Decoherence

I've skipped over an important topic in the preceding discussion. In order to model the interaction of environment and system, we really have two choices. The simple, rather brute-force method is to ignore the real nature of the interaction and model it as a measurement process. This is inherently an ad hoc process, however, unless we can actually develop a theory of the way that classical systems interact with quantum systems and cause them to collapse or whatever we decide that they do. The systematic study of this interaction – and, in particular, the study of how entangled states of the system and the environment fall apart into classically correlated but quantum mechanically independent states of system and environment is the **program of decoherence**. Unfortunately, this is far beyond our current scope, although the final example of a semiclassical dynamic process that I propose is closely related to decoherence processes.

6 Continuous Measurement: A Candidate Process

6.1 Background

Finally, after surveying that which has gone before, we can examine one form of semiclassical dynamics which has the potential to produce chaotic semiclassical behavior in phase space.

First, let's review the important characteristics of the sort of theory we're looking for. It should permit classical trajectories, and it should permit trajectories to diverge exponentially. Unless we give up entirely on correlations between Hilbert space and phase space, this means that it has to incorporate some nonunitary elements which do not preserve Hilbert space overlaps. In addition, we want to account for measurement processes as part of the dynamics, and the theory must connect smoothly with classical and quantum mechanics in the regimes in which each is demonstrably valid.

The idea of *continuous measurement* was apparently explored by Wojciech Zurek in the early 1990s, and has been more extensively developed by Michael Mensky of the Lebedev Institute in Moscow. For a much fuller description of the theory and its ramifications, the interested reader is referred to "Decoherence and the Theory of continuous quantum measurements" by M.B. Mensky in Physics Uspekhi 41, 923-940 (1998). I will present a brief intuitive derivation of the theory, due more to my own research than to Mensky's derivation.

6.2 Why Continuous Measurement?

We have seen previously that Gaussian wavepackets have particularly attractive properties in formulating semiclassical mechanics. A Gaussian wavepacket occupies a minimal volume of phase space (and is thus the nearest approach to the classical δ -function distribution of a point particle), is invariant under Fourier transforms, has a mean which coincides with its median and mode, and can be easily analyzed. Finally, a Gaussian wavepacket yields to a normal probability distribution, which fits measurement uncertainties well. For all these reasons, the theory of continuous measurement will be formulated in terms of Gaussian wavepackets.

As was also seen previously, wavepackets are semiclassically attractive – except that they spread. Any theory of semiclassical mechanics which relies on wavepackets must deal with the fact that they are only coherent in quadratic wells and over limited times as free particles. We would like

to find a way to make a wavepacket remain coherent independently of the external potential.

We begin with a generalization of position measurements. A measurement of \hat{x} on a state $\psi(x)$ that is consistent with the collapse postulate (Von Neumann reduction) yields two results: (1) an eigenvalue x_m of \hat{x} and (2) a final state $\psi_m(x)$ which is an eigenstate of \hat{x} . Motivated by need for measurements with uncertainty and measurements whose application does not destroy the prior state of the system (non-demolition measurements), we adopt the following measurement scheme.

6.2.1 Generalized Measurement with Uncertainty

A measurement is now defined not merely by the operator which it measures, but also by the precision to which it is measured. Thus a measurement of $\hat{x}_{\Delta x}$ yields three things: (1) a measured value x_m which need not necessarily be an eigenvalue of \hat{x} unless $\Delta x = 0$; (2) an uncertainty Δx , which we consider to have been predetermined by the nature of the measurement, and (3) a final state $\psi_m(x)$ given by the normalized product of the pre-measurement state with a *measurement envelope* determined by x_m , Δx , and the exact nature of the process represented by \hat{x} .

$$\psi_m(x) = A\psi(x) \cdot E_m(x_m; x) \quad (25)$$

where A is a normalizing constant, and an appropriate choice for the measurement envelope is

$$E_m(x_m; x) = \exp \left[-\frac{(x - x_m)^2}{4(\Delta x)^2} \right] \quad (26)$$

Whereas under the old theory the probability of a measurement yielding a value x_m was proportional to $|\psi(x_m)|^2$, under the new theory of measurement the probability of finding x_m is proportional to $\int dx E_m(x_m; x)|\psi(x)|^2$. This reflects the fact that even if $\psi(x_m) = 0$, because the measurement is unreliable we might measure x_m anyway. Obviously, if $E_m(x_m; x) = \delta(x_m - x)$, then we recover all properties of the old Von Neumann reduction.

Let's consider the properties of $\psi_m(x)$ as determined both by the initial state and the measurement process. If ψ has some overall modulation like $e^{\frac{i}{\hbar}px}$, then as long as the wavelength is smaller than the width of E_m , this modulation is largely undisturbed by the measurement process. The variance of the momentum distribution will increase somewhat, but this is a natural consequence of any x-measurement.

If the initial state is completely delocalized, then the shape of ψ_m is determined almost entirely by E_m , and the resulting probability distribution $P(x) = |\psi^2(x)|$ is a Gaussian distribution centered around $\langle x \rangle = x_m$, with variance Δx – exactly what we would expect after a measurement $\hat{x} = x_m \pm \Delta x$. Thus if the initial state cannot possibly give us any information about \hat{x} , then the final state is consistent with the measurement. This was the original intention of Von Neumann reduction (the collapse postulate); the final state had to be consistent with the measurement record.

If the initial state is very nearly a delta function of position, then the normalized product $\psi \times E_m$ will be dominated by ψ . In this case, if we had the maximum knowledge of the initial state possible we would conclude that the result of our measurement was insignificant compared to our already-precise knowledge of the prior state of the particle, and so the final state should reflect the initial state more than the outcome of the measurement. Again, the final state is consistent with this interpretation.

In general, if the initial state is a Gaussian wavepacket centered at x_0 with variance Δx_0 , then the final state can be calculated explicitly. The final state (in the \hat{x} representation) is also a Gaussian wavepacket with the properties:

$$x_f = \frac{\Delta x_0 x_m + \Delta x_m x_0}{(\Delta x_0)^2 + (\Delta x_m)^2} \quad (27)$$

$$\Delta x_f = \frac{\Delta x_0 \Delta x_m}{\sqrt{(\Delta x_0)^2 + (\Delta x_m)^2}} \quad (28)$$

If the initial state had a complex phase, then the final state also has a complex phase, but the size of the perturbation in momentum space due to the measurement is a non-trivial function of the measured position and the measurement strength.

We can see from these relations that all the relationships which we expect from experimental error and uncertainty hold for this measurement scheme.

6.2.2 Periodic Measurement Dynamics

As is apparent in equation 28, the effect of a measurement is always to reduce the width of the state. We know that wavepackets spread out due to unitary time evolution, but if we suppose that the system is in contact with an environment which periodically performs a measurement on it, then we can model this interaction as an alternating series of unitary time evolutions and weak measurements of the type just outlined. Thus what we have postulated is a time evolution operator \hat{V} which is built out of unitary evolutions over

a time Δt_α ($\hat{U}_{\Delta t}$) and measurements of position ($\hat{X}_{x_\alpha, \Delta x_\alpha}$). If we make the simplifying assumption that the measurements are periodic instead of stochastic and that each measurement is equally precise, then \hat{V} looks like this:

$$\begin{aligned}\hat{V}(n\Delta t) &= \prod_{i=1}^n \hat{X}_{x_i, \Delta x} \hat{U}_{\Delta t} \\ &= \hat{X}_{x_n, \Delta x} \hat{U}_{\Delta t} \hat{X}_{x_{n-1}, \Delta x} \dots \hat{U}_{\Delta t} \hat{X}_{x_1, \Delta x} \hat{U}_{\Delta t}\end{aligned}\quad (29)$$

This is a propagation operator along some path in coordinate space determined by the $\{x_i\}$'s, which can be broken up into many tiny steps along the path. The path-integral formulation immediately comes to mind, and Mensky has derived the properties of continuous measurement from a restricted path integral determined by the sequence of measurements. We can go further than this, however, because we can suppose a similar and simultaneous process occurring in momentum space – that is, a series of periodic weak measurements of position. We can describe such an evolution operator with alternating position and momentum measurements:

$$\hat{V} \approx \hat{X}_{x_n, \Delta x} \hat{U}_{\frac{1}{2}\Delta t} \hat{P}_{p_n, \Delta p} \hat{U}_{\Delta t} \hat{X}_{x_{n-1}, \Delta x} \dots \quad (30)$$

Since each measurement localizes its corresponding operator, such an evolution operator would tend to maintain localization in both position and momentum while obeying Ehrenfest's theorem for sufficiently weak measurements.

6.2.3 The Effects of Continuous Measurement

Mensky's path-integral formulation requires a measurement record around which path integrals are calculated. We can take this idea further, however, by considering what happens if the accuracy of each measurement and the time between successive measurements are simultaneously shrunk, so that we take a limit as $\Delta x, \Delta p \rightarrow \infty$, $\Delta t \rightarrow 0$. Then we can consider a long string of very weak measurements, ignoring the time evolution in between them. The product of a series of N $\hat{X}_{x_i, \Delta x}$ operators is a multiplication by the product of all the measurement envelopes $E_m(x_i; x)$:

$$\left(\prod_N \hat{X}_{x_i, \Delta x} \right) \psi(x) = \left(\prod_N E_m(x_i; x) \right) \psi(x) \quad (31)$$

$$= e^{-\frac{(x-x_1)^2}{4\Delta x^2}} e^{-\frac{(x-x_2)^2}{4\Delta x^2}} e^{-\frac{(x-x_3)^2}{4\Delta x^2}} \dots \quad (32)$$

$$= \exp \left[-\frac{1}{4\Delta x^2} \left[(x - x_1)^2 + (x - x_2)^2 + \dots \right] \right] \quad (33)$$

$$= \exp \left[-\frac{N}{4\Delta x^2} (x - \bar{x}_i)^2 \right] \quad (34)$$

Since all the x_i 's are probabilistic, their mean is just the expectation value of \hat{x} , $\langle x \rangle$, in the limit that there are infinitely many x_i 's measured in each amount of time. We conclude that if the measurement process is continuous, then all the measurements average to a single effective measurement centered at $\langle x \rangle$. The consequences of this conclusion are staggering.

First, we have succeeded in defining a plausible time evolution scheme which (1) is deterministic (like classical mechanics), (2) involves a measurement process, (3) provides a measurement record, (4) maintains a localized wavepacket in both position and momentum for indefinite time, and (5) can be shown to reduce, in the limit of weak or infrequent measurements and in the limit of low mass, to the unitary evolution that we observe in the highly quantum regime. Second, this scheme has an elegant physical interpretation. A particle in a bath of low-energy photons, which is constantly scattering the photons, necessarily becomes entangled (correlated) with the photons it scatters. If a decoherence mechanism can be demonstrated by which the entanglement of the photons and the particle is exchanged for classical correlations (this, by the way, is exactly the process for which decoherence was invented), then that represents a measurement, and this scheme is realized.

Finally, because the measurements reduce to a deterministic process in the limit of frequent measurements, we can actually write this evolution process as a special case of the unitary evolution operator $\hat{U} = \exp \left[-\frac{i}{\hbar} \hat{H} t \right]$, if we add an imaginary term to the Hamiltonian to represent the measurement process. Complex Hamiltonians have been used to describe particle creation and destruction; in this case we can think of them as representing the reshaping of a particle by selective local construction and destruction. The necessary Hamiltonian (presented without proof) is:

$$\hat{H}_{\text{measurement}} = -i \frac{\alpha}{(\Delta x)^2} \left((x - \langle x \rangle)^2 - 2 \right) \quad (35)$$

The parameter α is dimensionless and controls the strength of the measurement; it's related to the frequency of measurement. The $(x - \langle x \rangle)^2$ term represents the averaged out measurement envelope. The odd-looking -2 stuck on the end really is supposed to be there; it maintains normalization, and comes originally from the constant A in equation 25.

This result opens far more doors than it closes – if nothing else, it would be nice to do an experiment to see if something like this actually seems to be

happening. In addition, many more consequences of this dynamics need to be worked out. For now, though, I will close with two comments about the one desirable criteria for a theory of semiclassical dynamics that I haven't already shown this model to satisfy, and another comment on symmetry properties. Recall that we required that the theory produce exponential divergence of trajectories and involve some sort of nonlinearity in order to permit information production. Clearly, since this model produces stable wavepackets that follow classical trajectories, it satisfies the first criteria. Its interaction with the environment permits information production by, in a sense, distilling information from the environment. But where is the nonlinearity? The answer is: the measurement Hamiltonian is inherently nonlinear because it involves a term, $\langle x \rangle$, which is dependent on the wavefunction. This permits the sort of exponential behavior that we examined in the section on chaotic behavior.

Finally, I would like to point out that with this model we have rather accidentally recovered another aspect of classical mechanics. Classical Hamiltonian mechanics is formulated symmetrically in position and momentum; presumably a semiclassical theory should demonstrate some similar trait. If we recall that the measurement process (which completely determines the imaginary Hamiltonian) can be formulated in a symmetric way, then it is apparent that the full measurement Hamiltonian due to measurements of both position and momentum is also symmetric in x and p . We cannot really interpret this as any kind of classical Hamiltonian, however; classical mechanics is formulated with real numbers. Thus the measurement "Hamiltonian," which is completely independent of the real Hamiltonian, must actually represent a feature of semiclassical dynamics *in general*, and not a system-dependent feature at all. Thus, our semiclassical formulation of mechanics is not only symmetric in position and momentum, but is (as formulated) quadratic in both – a very natural state of affairs.

7 Conclusions

I hope that I have succeeded in demonstrating both some of the problems and some of the solutions in semiclassical chaos. We have examined aspects of quantum mechanics, of classical chaos, and of pure classical dynamics. The challenge of this field is to bring all of these together to form a coherent whole. Unfortunately, the amount of material covered here has precluded my discussing a few other aspects of this work which are of at least academic interest; in particular, the connection to astrophysical mass distributions has

been left out entirely. To those who chose to read this paper expecting to find astrophysics – my sincere apologies.

I would particularly like to thank a few people for their help in one way or another. My sincere gratitude goes to Ben Schumacher for getting it all started, and to Robert Littlejohn for consultation and also for writing an extraordinarily clear review article which formed the base of much of this paper.

8 Bibliography

- Littlejohn, Robert. "The Semiclassical evolution of Wave Packets," in Physics Reports, v. 138 n. 4 (May, 1986)
- Mensky, M.B. "Decoherence and the theory of continuous quantum measurements," in Physics Uspekhi 41, 923-940 (1998)
- Balazs, N.L. and Jennings, B.K. "Wigner's function and other distribution functions in mock phase spaces" in Physics Reports v. 104 n. 6 (Feb, 1984)
- Giulini, D. et al. *Decoherence and the Appearance of a Classical World in Quantum Theory*, Springer: 1996.
- Gutzwiller, M.C. *Chaos in Classical and Quantum Mechanics*, Springer-Verlag, 1990.

A Classical Distributions

Classical particles are usually assumed to exist at a particular coordinate \vec{q} . This is simply the definition of a “point particle.” Under three circumstances, however, we may find it useful to describe a classical particle by a probabilistic distribution in phase space⁷

1. Time-dependent forces such as Brownian motion may prevent us from knowing the exact state of the particle at any time. This, however, implies a time-dependent Hamiltonian, and requires statistical mechanics. Such distributions are often governed by stochastic evolution equations (the Fokker-Planck equation, for instance), and are outside the scope of this paper.
2. We may simply be uncertain of the position of the particle. This is always the case in practice; no measurement apparatus can make a perfect measurement. The most common distribution function encountered in this case is an asymmetric Gaussian (asymmetric in position and momentum), although any distribution is conceivable. We implicitly assumed this case in the last section when discussing the set S of points within a phase space volume \mathcal{V} .
3. Finally, we may actually be interested in the motion of a set of particles. This is often the case in astrophysics, where dust clouds, masses of gas, and star clusters obey distribution laws. In this case, however, we must be extremely careful to discriminate between three very distinct cases. The particles may be truly non-interacting (dust clouds are probably the best example), in which case they form a Liouville distribution. The particles may have interactions which are inherently stochastic or statistical in nature (gas clouds, for instance) in which case we’re really back to case 1 and the Fokker-Planck equation. Finally, the particles may interact according to well-defined Hamiltonian interactions, which leads to the general (and extremely interesting) case of self-interacting distributions, which I discuss later. Particular examples of this case are globular clusters and charged plasmas.

The most easily treated of these cases is the case of non-interacting, non-stochastic particles. In this case, we have what is called a Liouville dis-

⁷Generally, we think of a distribution in coordinate (\vec{q}) space; in general, however, the distribution is over all $2N$ dimensions of phase space. A coordinate-only distribution would just be the product of an N -dimensional \vec{q} -distribution with an N -dimensional $\delta(\vec{p} - \vec{p}_0)$ function.

tribution – a classical “swarm of particles”. The time evolution of Liouville distributions is governed by the linear equation:

$$\frac{\partial f(\vec{q}, \vec{p})}{\partial t} = -[f(\vec{q}, \vec{p}), H] \quad (36)$$

where the *Poisson Bracket* $[f, H]$ is defined as:

$$[f, H]_{\vec{q}, \vec{p}} = \vec{\nabla}_q f \vec{\nabla}_p H - \vec{\nabla}_p f \vec{\nabla}_q H \quad (37)$$

This formalism permits us to make a few sweeping statements about the dynamics of distributions. First of all, distributions in phase space move like an incompressible fluid.⁸ The Liouville equation is linear, and does not admit chaotic solutions; the average values of dynamical variables $\langle \vec{q} \rangle$ and $\langle \vec{p} \rangle$ evolve linearly in time. Finally, if we define the *overlap* of two normalized distributions as:

$$P(f_1, f_2) = \int d^N q d^N p (f_1(\vec{q}, \vec{p}) f_2(\vec{q}, \vec{p}))^{\frac{1}{2}} \quad (38)$$

then the overlap of two distributions f_1 and f_2 is constant in time. This is, again, apparent from Liouville’s theorem, since any differential element of phase space volume $d^{2N}z$ which contains some element of each distribution df_1, df_2 at one time must evolve to a new element with the same volume $d^{2N}z$ which contains the same elements of the distributions.

We use these concepts to compare classical trajectories with quantum wavepacket motions, but for now let’s take a moment to consider some quantum mechanical issues.

⁸The incompressibility of this phase space “fluid” is simply a statement of Liouville’s theorem concerning volume invariance, which is in turn a consequence of a theorem by Poincare concerning the invariance of phase space volume under canonical transformations.

B Delocalization and Phase Space in Quantum Mechanics

One of the most striking differences between classical and quantum mechanics, which must be dealt with in order to draw connections between the two, is delocalization. I don't mean by this the question of locality in quantum theories, but the fact that **all** systems in quantum theory are spread out in phase space. Even though we haven't defined quantum phase space yet, the uncertainty principle is familiar to any physicist. If we choose to express a quantum state in a position (\hat{q}) or momentum (\hat{p}) basis, then any attempt that we make to localize the state in one basis increases its delocalization, or *spread*, in the other basis. The minimal simultaneous uncertainty in x and p is given by the Heisenberg relation:

$$\Delta x \Delta p \geq \frac{\hbar}{2} \quad (39)$$

where the state which satisfies the corresponding equality (and thus is maximally localized) is of course a Gaussian wavepacket. The Gaussian wavepacket is the first eigenstate of the simple harmonic oscillator, which has the unique property of being the only quadratic Hamiltonian which is symmetric in position and momentum.⁹ The minimum phase space volume which a quantum state in an N -dimensional coordinate space can occupy is thus $\left(\frac{\hbar}{2}\right)^N$.

Interestingly enough, the phase space volume occupied by a quantum state decreases exponentially with N . We have already noted a similar phenomenon in Hilbert space; as the number of independent dimensions of Hilbert space goes up, the overlap between states separated by a small angle θ in each of the N independent dimensions of Hilbert space decreases as $\cos(\theta)^N$. This indicates the possibility of drawing connections between classical and quantum phase space, so why don't we explore that topic before going on to consider classical delocalization.

⁹Which explains why its eigenstate is maximally localized in both position and momentum. Interestingly enough, you can construct a *linear* Hamiltonian which is symmetric in position and momentum, $H = x + p = -i\hbar \frac{\partial}{\partial x} + x$, whose eigenfunctions are of the form $\psi(x) = \exp[-\frac{i}{2\hbar}x^2]$ and have in some sense *minimal* simultaneous localization, since $\psi^2(x) = \psi^2(p) = \text{constant}$. This is of no particular use, since all physical Hamiltonians are at least quadratic in momentum (the energy eigenvalues of the linear symmetric Hamiltonian are all the complex numbers, which makes it quite obvious that it's a nonphysical Hamiltonian)

B.1 Quantum Mechanics in Phase Space: The Wigner Distribution

The Wigner distribution has been around since the early days of quantum mechanics; Wigner essentially asked himself, “What is the quantum mechanical equivalent of phase space?” and came up with the distribution that bears his name. The formal theory behind the Wigner distribution (or Wigner function) is that of *symbols* and *Weyl associations*. A symbol is, in general, a way of connecting an operator on a Hilbert space with a function or functions defined on some space of classical variables. For instance, we can interpret the matrix element of an operator \hat{A} , $\langle \vec{x} | \hat{A} | \vec{x}' \rangle$ as a symbol on the $2N$ dimensional classical space (\vec{x}, \vec{x}') . The Weyl symbol for a given operator is a particularly convenient symbol defined by:

$$a_W(\vec{u}, \vec{v}) = \text{Tr}[\hat{T}(\vec{u}, \vec{v})\hat{A}] \quad (40)$$

where \vec{u} and \vec{v} are classical variables and the \hat{T} operator is defined by:

$$\hat{T}(\vec{u}, \vec{v}) = \exp \left[i(\vec{u} \cdot \hat{\vec{p}} + \vec{v} \cdot \hat{\vec{q}}) \right] \quad (41)$$

For our purposes, this is all deep background; the important thing is that the Wigner distribution is a specific Weyl-type symbol for the density operator, defined as:

$$f_W(\vec{p}, \vec{q}) = \text{Tr}(\hat{\Delta}\hat{\rho}) \quad (42)$$

where the $\hat{\Delta}$ operator is essentially a transformation of the \hat{T} operator from (\vec{u}, \vec{v}) space to (\vec{q}, \vec{p}) space:

$$\hat{\Delta}(\vec{p}', \vec{q}') = \frac{\hbar^3}{(2\pi)^3} \int \exp \left[i \left(\begin{array}{c} \vec{u} \cdot (\hat{\vec{p}} - \vec{p}') \\ + \\ \vec{v} \cdot (\hat{\vec{q}} - \vec{q}') \end{array} \right) \right] d^N u d^N v \quad (43)$$

Since the general intent of a function is to provide a link between a quantum operator and some classical function, and since the density operator deals with probability densities for a quantum state, it seems sensible that the symbol associated with the density operator should be a semi-classical probability density function – that is, a phase space distribution. This is exactly the interpretation of the Wigner function, which can more practically be written:

$$f_W(\vec{p}, \vec{q}) = \int d^N s \exp \left(\frac{i}{\hbar} \vec{s} \cdot \vec{p} \right) \rho \left(\vec{q} - \frac{1}{2} \vec{s}, \vec{q} + \frac{1}{2} \vec{s} \right) \quad (44)$$

for a system with N independent coordinates. If we rewrite $\hat{\rho} = |\psi\rangle\langle\psi|$ and note that in the preceding equation, $\rho\left(\vec{q} - \frac{1}{2}\vec{s}, \vec{q} + \frac{1}{2}\vec{s}\right) = \left\langle \vec{q} + \frac{1}{2}\vec{s} \left| \hat{\rho} \right| \vec{q} - \frac{1}{2}\vec{s} \right\rangle$, then we arrive at the following (and this is the last one!) formulation of the Wigner distribution:

$$f_W(\vec{p}, \vec{q}) = \int d^N s \exp\left(\frac{i}{\hbar}\vec{s} \cdot \vec{p}\right) \psi\left(\vec{q} + \frac{1}{2}\vec{s}\right) \psi^*\left(\vec{q} - \frac{1}{2}\vec{s}\right) \quad (45)$$

This last formula is less general than the previous formulation because, in utilizing the definition of the density matrix in terms of a single state, we limit ourselves to pure states. However, by eliminating the density matrix, we save a great deal of complication.

Having so exhaustively defined the Wigner function, we might ask what good it is. I shall not attempt to enumerate all the properties of the Wigner function here, nor to actually apply it to a physical system, but merely discuss some of the highlights. The Wigner function acts in most ways exactly like a classical Liouville distribution in $2N$ -dimensional phase space. Integrating over all values of \vec{p} yields the probability distribution $\psi^2(\vec{q})$, and integrating over the coordinate yields the appropriate momentum-space distribution. The Wigner function is real-valued, as a probability distribution should be, and if we integrate along any line in phase space (to extract the probability of the system in a set of states $\vec{p} = \vec{P}(\vec{q})$, the result is always positive. However (and this is the major difference between the Wigner function and a Liouville distribution), the Wigner function **may** take on negative values. This simply emphasizes the fact that quantum mechanics does not permit simultaneous localization in \vec{p} and \vec{q} ; if we try to evaluate $f_W(\vec{q}, \vec{p})$ at a point, we get ridiculous results. Any operation on f_W , however, which is physically realizable (i.e., an integration of f_W over a volume of phase space greater than \hbar^N) will yield a positive, physically sensible result.

The wave function holds all available quantum information (for a pure state) about both position and momentum. All of this information is transferred to the Wigner distribution¹⁰ *except* the overall phase information. In addition, the Wigner function is a nonlinear function of $|\psi\rangle$ – that is, it doesn't obey superposition. We can conclude that the Wigner function displays many of the properties we associate with classical (as opposed to

¹⁰Actually, the Wigner function is more general than the wavefunction; not only is there a well-defined mapping from pure state wavefunctions to Wigner functions, but also a 1:1 mapping from density matrices to Wigner functions. In addition, each Wigner function is a real-valued, normalized function on phase space, but not every real-valued and normalized function on phase space is a valid Wigner function; the problem of resolving whether a given $f(\vec{q}, \vec{p})$ is a valid Wigner function is not generally resolved.

quantum) distributions; it is nonlinear (which means that Wigner function overlaps need not be preserved under unitary evolution) and a wavepacket which is separated into two uncorrelated wavepackets and recombined will not display interference because Wigner functions have no overall phase information. This does not mean that Wigner functions have no phase information; if we write a single Wigner function for a wavepacket split into two parts, then the relative phase of the two parts remains intact and can be recovered from the Wigner function. If, however, the wavepacket is completely sundered in some fashion (an irreversible process), then we can consider each of the two parts to be an independent state, and describe each by a separate Wigner function, in which case relative phase (which is now the overall phase of each wave packet) is lost.