

Lectures on the Geometry of Quantization

Sean Bates

Department of Mathematics

Columbia University

New York, NY 10027 USA

`smb@math.columbia.edu`

Alan Weinstein

Department of Mathematics

University of California

Berkeley, CA 94720 USA

`alanw@math.berkeley.edu`

Contents

1	Introduction: The Harmonic Oscillator	5
2	The WKB Method	8
2.1	Some Hamilton-Jacobi preliminaries	8
2.2	The WKB approximation	11
3	Symplectic Manifolds	17
3.1	Symplectic structures	17
3.2	Cotangent bundles	28
3.3	Mechanics on manifolds	32
4	Quantization in Cotangent Bundles	36
4.1	Prequantization	36
4.2	The Maslov correction	41
4.3	Phase functions and lagrangian submanifolds	46
4.4	WKB quantization	56
5	The Symplectic Category	64
5.1	Symplectic reduction	64
5.2	The symplectic category	76
5.3	Symplectic manifolds and mechanics	79
6	Fourier Integral Operators	83
6.1	Compositions of semi-classical states	83
6.2	WKB quantization and compositions	87
7	Geometric Quantization	93
7.1	Prequantization	93
7.2	Polarizations and the metaplectic correction	98
7.3	Quantization of semi-classical states	109
8	Algebraic Quantization	111
8.1	Poisson algebras and Poisson manifolds	111
8.2	Deformation quantization	112
8.3	Symplectic groupoids	114
A	Densities	119
B	The method of stationary phase	121
C	Čech cohomology	124
D	Principal \mathbb{T}_\hbar bundles	125

Preface

These notes are based on a course entitled “Symplectic geometry and geometric quantization” taught by Alan Weinstein at the University of California, Berkeley, in the fall semester of 1992 and again at the Centre Emile Borel (Institut Henri Poincaré) in the spring semester of 1994. The only prerequisite for the course (and for these notes) was a knowledge of the basic notions from the theory of differentiable manifolds (differential forms, vector fields, transversality, etc.). The aim of the course was to give students an introduction to the ideas of microlocal analysis and the related symplectic geometry, with an emphasis on the role which these ideas play in formalizing the transition between the mathematics of classical dynamics (hamiltonian flows on symplectic manifolds) and that of quantum mechanics (unitary flows on Hilbert spaces).

There already exist many books on the subjects treated here, but most of them provide too much detail for the reader who just wants to find out what the subject is about. These notes are meant to function as a guide to the literature; we refer to other sources for many details which are omitted here, and which can be bypassed on a first reading.

The pamphlet [63] is in some sense a precursor to these notes. On the other hand, a much more complete reference on the subject, written at about the same time, is [28]. An earlier work, one of the first to treat the connections between classical and quantum mechanics from a geometric viewpoint, is [41]. The book [29] treats further topics in symplectic geometry and mechanics, with special attention to the role of symmetry groups, a topic pretty much ignored in the present notes. For more extensive treatment of the PDE aspects of the subject, we refer to [43] for a physics-oriented presentation and to the notes [21] and the treatises [32], [46], and [56]. For “geometric quantization”, one may consult [35], [53], [54], [60] or [71]. For classical mechanics and symplectic geometry, we suggest [1], [2], [6], [8], [25], [38], [59]. Finally, two basic references on quantum mechanics itself are [13] and [20].

Although symplectic geometry is like any field of mathematics in having its definitions, theorems, etc., it is also a special way of looking at a very broad part of mathematics and its applications. For many “symplecticians”, it is almost a religion. A previous paper by one of us [64] referred to “the symplectic creed”.¹ In these notes, we show how symplectic geometry arises from the study of semi-classical solutions to the Schrödinger equation, and in turn provides a geometric foundation for the further analysis of this and other formulations of quantum mechanics.

These notes are still not in final form, but they have already benefitted from the comments

¹We like the following quotation from [4] very much:

In recent years, symplectic and contact geometries have encroached on all areas of mathematics. As each skylark must display its comb, so every branch of mathematics must finally display symplectisation. In mathematics there exist operations on different levels: functions acting on numbers, operators acting on functions, functors acting on operators, and so on. Symplectisation belongs to the small set of highest level operations, acting not on details (functions, operators, functors), but on all the mathematics at once. Although some such highest level operations are presently known (for example, algebraisation, Bourbakisation, complexification, superisation, symplectisation) there is as yet no axiomatic theory describing them.

and suggestions of many readers, especially Maurice Garay, Jim Morehead, and Dmitry Roytenberg. We welcome further comments. We would like to thank the Centre Emile Borel and the Isaac Newton Institute for their hospitality. During the preparation of these notes, S.B. was supported by NSF graduate and postdoctoral fellowships in mathematics. A.W. was partially supported by NSF Grants DMS-90-01089 and 93-01089.

1 Introduction: The Harmonic Oscillator

In these notes, we will take a “spiral” approach toward the quantization problem, beginning with a very concrete example and its proposed solution, and then returning to the same kind of problem at progressively higher levels of generality. Specifically, we will start with the harmonic oscillator as described classically in the phase plane \mathbb{R}^2 and work toward the problem of quantizing arbitrary symplectic manifolds. The latter problem has taken on a new interest in view of recent work by Witten and others in the area of topological quantum field theory (see for example [7]).

The classical picture

The harmonic oscillator in 1 dimension is described by Newton’s differential equation:

$$m\ddot{x} = -kx.$$

By a standard procedure, we can convert this second-order ordinary differential equation into a system of two first-order equations. Introducing the “phase plane” \mathbb{R}^2 with position and momentum coordinates (q, p) , we set

$$q = x \quad p = m\dot{x},$$

so that Newton’s equation becomes the pair of equations:

$$\dot{q} = \frac{p}{m} \quad \dot{p} = -kq.$$

If we now introduce the **hamiltonian function** $H: \mathbb{R}^2 \rightarrow \mathbb{R}$ representing the sum of kinetic and potential energies,

$$H(q, p) = \frac{p^2}{2m} + \frac{kq^2}{2}$$

then we find

$$\dot{q} = \frac{\partial H}{\partial p} \quad \dot{p} = -\frac{\partial H}{\partial q}$$

These simple equations, which can describe a wide variety of classical mechanical systems with appropriate choices of the function H , are called **Hamilton’s equations**.² Hamilton’s equations define a flow on the phase plane representing the time-evolution of the classical system at hand; solution curves in the case of the harmonic oscillator are ellipses centered at the origin, and points in the phase plane move clockwise around each ellipse.

We note two qualitative features of the hamiltonian description of a system:

1. The derivative of H along a solution curve is

$$\frac{dH}{dt} = \frac{\partial H}{\partial q} \dot{q} + \frac{\partial H}{\partial p} \dot{p} = -\dot{p}\dot{q} + \dot{q}\dot{p} = 0,$$

²If we had chosen \dot{x} rather than $m\dot{x}$ as the second coordinate of our phase plane, we would not have arrived at this universal form of the equations.

i.e., the value of H is constant along integral curves of the hamiltonian vector field. Since H represents the total energy of the system, this property of the flow is interpreted as the law of conservation of energy.

2. The divergence of the hamiltonian vector field $X_H = (\dot{q}, \dot{p}) = (\frac{\partial H}{\partial p}, -\frac{\partial H}{\partial q})$ is

$$\nabla \cdot X_H = \frac{\partial^2 H}{\partial q \partial p} - \frac{\partial^2 H}{\partial p \partial q} = 0.$$

Thus the vector field X_H is divergence-free, and its flow preserves area in the phase plane.

The description of classical hamiltonian mechanics just given is tied to a particular coordinate system. We shall see in Chapter 3 that the use of differential forms leads to a coordinate-free description and generalization of the hamiltonian viewpoint in the context of symplectic geometry.

The quantum mechanical picture

In quantum mechanics, the motion of the harmonic oscillator is described by a complex-valued wave function $\psi(x, t)$ satisfying the 1-dimensional **Schrödinger equation**:

$$i\hbar \frac{\partial \psi}{\partial t} = -\frac{\hbar^2}{2m} \frac{\partial^2 \psi}{\partial x^2} + \frac{k}{2} x^2 \psi.$$

Here, *Planck's constant* \hbar has the dimensions of action (energy \times time). Interpreting the right hand side of this equation as the result of applying to the wave function ψ the operator

$$\hat{H} \stackrel{def}{=} -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + \frac{k}{2} m_{x^2},$$

where m_{x^2} is the operator of multiplication by x^2 , we may rewrite the Schrödinger equation as

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

A solution ψ of this equation does not represent a classical trajectory; instead, if ψ is *normalized*, i.e.

$$\int_{\mathbb{R}} \psi^* \psi = 1,$$

then its square-norm

$$\rho(x, t) = |\psi(x, t)|^2$$

is interpreted as a probability density for observing the oscillator at the position x at time t . The wave function $\psi(x, t)$ itself may be viewed alternatively as a t -dependent function of x , or as a path in the function space $C^\infty(\mathbb{R}, \mathbb{C})$. From the latter point of view, Schrödinger's equation defines a vector field on $C^\infty(\mathbb{R}, \mathbb{C})$ representing the time evolution of the quantum system; a wave function satisfying Schrödinger's equation then corresponds to an integral curve of the associated flow.

Like Hamilton's equations in classical mechanics, the Schrödinger equation is a general form for the quantum mechanical description of a large class of systems.

Quantization and the classical limit

The central aim of these notes is to give a geometric interpretation of relationships between the fundamental equations of classical and quantum mechanics. Based on the present discussion of the harmonic oscillator, one tenuous connection can be drawn as follows. To the classical position and momentum observables q, p we associate the differential operators

$$q \mapsto \hat{q} = m_x$$
$$p \mapsto \hat{p} = -i\hbar \frac{\partial}{\partial x}.$$

The classical hamiltonian $H(q, p) = p^2/2m + kq^2/2$ then corresponds naturally to the operator \hat{H} . As soon as we wish to “quantize” a more complicated energy function, such as $(1 + q^2)p^2$, we run in to the problem that the operators \hat{q} and \hat{p} do not commute with one another, so that we are forced to choose between $(1 + \hat{q}^2)\hat{p}^2$ and $\hat{p}^2(1 + \hat{q}^2)$, among a number of other possibilities. The difference between these choices turns out to become small when $\hbar \rightarrow 0$. But how can a constant approach zero?

Besides the problem of “quantization of equations,” we will also treat that of “quantization of solutions.” That is, we would like to establish that, for systems which are in some sense macroscopic, the classical motions described by solutions of Hamilton’s equations lead to approximate solutions of Schrödinger’s equation. Establishing this relation between classical and quantum mechanics is important, not only in verifying that the theories are consistent with the fact that we “see” classical behavior in systems which are “really” governed by quantum mechanics, but also as a tool for developing approximate solutions to the quantum equations of motion.

What is the meaning of “macroscopic” in mathematical terms? It turns out that good approximate solutions of Schrödinger’s equation can be generated from classical information when \hbar is small enough. But how can a constant with physical dimensions be small?

Although there remain some unsettled issues connected with the question, “How can \hbar become small?” the answer is essentially the following. For any particular mechanical system, there are usually characteristic distances, masses, velocities, . . . from which a unit of action appropriate to the system can be derived, and the classical limit is applicable when \hbar divided by this unit is much less than 1. In these notes, we will often regard \hbar mathematically as a formal parameter or a variable rather than as a fixed number.

2 The WKB Method

A basic technique for obtaining approximate solutions to the Schrödinger equation from classical motions is called the *WKB method*, after Wentzel, Kramers, and Brillouin. (Other names, including Liouville, Green, and Jeffreys are sometimes attached to this method. References [13] and [47] contain a discussion of some of its history. Also see [5], where the method is traced all the way back to 1817. For convenience, nevertheless, we will still refer to the method as WKB.) A good part of what is now called microlocal analysis can be understood as the extension of the basic WKB idea to more precise approximations and more general situations, so the following discussion is absolutely central to these notes.

2.1 Some Hamilton-Jacobi preliminaries

In this section we will carry out the first step in the WKB method to obtain an approximate solution to the “stationary state” eigenvalue problem arising from the Schrödinger equation. The geometric interpretation of this technique will lead to a correspondence between classical and quantum mechanics which goes beyond the one described in Chapter 1.

Consider a 1-dimensional system with hamiltonian

$$H(q, p) = \frac{p^2}{2m} + V(q),$$

where $V(x)$ is a potential (for example the potential $kx^2/2$ for the harmonic oscillator). Hamilton’s equations now become

$$\dot{q} = \frac{p}{m} \quad \dot{p} = -V'(q).$$

For fixed $\hbar \in \mathbb{R}_+$, Schrödinger’s equation assumes the form

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

where

$$\hat{H} = -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + m_V$$

is the **Schrödinger operator**.

As a first step toward solving the Schrödinger equation, we look for **stationary states**, i.e. solutions of the form

$$\psi(x, t) = \varphi(x) e^{-i\omega t},$$

so-called because as time evolves, these solutions keep the same form (up to multiplication by a complex scalar of norm 1). Substituting this expression for ψ in the Schrödinger equation, we obtain

$$\hbar\omega \varphi(x) e^{-i\omega t} = (\hat{H}\varphi)(x) e^{-i\omega t}.$$

Eliminating the factor $e^{-i\omega t}$ above, we arrive at the **time-independent Schrödinger equation**:

$$(\hat{H} - E) \varphi = 0,$$

where $E = \hbar\omega$. This equation means that φ is to be an eigenfunction of the linear differential operator \hat{H} ; the eigenvalue E represents the energy of the system, which has a “definite value” in this state.

Suppose for the moment that the potential V is constant, in which case the force $-V'(x)$ is zero, and so we are dealing with a *free particle*. Trying a solution of the form $\varphi(x) = e^{ix\xi}$ for some constant ξ , we find that

$$(\hat{H} - E)\varphi = 0 \quad \Leftrightarrow \quad (\hbar\xi)^2 = 2m(E - V).$$

For $V < E$, the (real) value of ξ is thus determined up to a choice of sign, and one has an abundance of exact solutions of the Schrödinger equation which are oscillatory and bounded. Such a wave function is not square-integrable and as such is said to be “unnormalizable”; it represents a particle which is equally likely to be anywhere in space, but which has a definite momentum (since it is an eigenfunction of the momentum operator \hat{p}).³ When $E < V$, the constant ξ is imaginary, and there are only real exponential solutions, which are unbounded and admit no physical interpretation.

The basic idea at this stage of the WKB method is that, if V varies with x , then ξ should vary with x as well; a more general solution candidate is then

$$\varphi(x) = e^{iS(x)/\hbar},$$

for some real-valued function S known as a *phase function*. This proposed form of the solution is the simplest version of the *WKB ansatz*, and in this case we have

$$(\hat{H} - E)\varphi = \left[\frac{(S'(x))^2}{2m} + (V - E) - \frac{i\hbar}{2m}S''(x) \right] e^{iS(x)/\hbar}.$$

Since we will consider \hbar to be “small”, our first-order approximation attempt will ignore the last term in brackets; to kill the other two terms, we require that S satisfy the **eikonal** or **Hamilton-Jacobi equation**:

$$H(x, S'(x)) = \frac{(S'(x))^2}{2m} + V(x) = E,$$

i.e.

$$S'(x) = \pm\sqrt{2m(E - V(x))}.$$

To understand the phase function S geometrically, we consider the classical phase⁴ plane $\mathbb{R}^2 \simeq T^*\mathbb{R}$ with coordinates (q, p) . The differential $dS = S' dx$ can be viewed as a mapping $dS: \mathbb{R} \rightarrow T^*\mathbb{R}$, where as usual we set $p = S'$. Then S satisfies the Hamilton-Jacobi equation if and only if the image of dS lies in the level manifold $H^{-1}(E)$. This observation establishes a fundamental link between classical and quantum mechanics:

When the image of dS lies in a level manifold of the classical hamiltonian, the function S may be taken as the phase function of a first-order approximate solution of Schrödinger’s equation.

³See [55] for a group-theoretic interpretation of such states.

⁴These two uses of the term “phase” seem to be unrelated!

The preceding discussion generalizes easily to higher dimensions. In \mathbb{R}^n , the Schrödinger operator corresponding to the classical hamiltonian

$$H(q, p) = \frac{\sum p_i^2}{2m} + V(q)$$

is

$$\hat{H} = -\frac{\hbar^2}{2m}\Delta + mV,$$

where Δ denotes the ordinary Laplacian. As before, if we consider a WKB ansatz of the form $\varphi = e^{iS/\hbar}$, then

$$(\hat{H} - E)\varphi = \left[\frac{\|\nabla S\|^2}{2m} + (V - E) - \frac{i\hbar}{2m}\Delta S \right] e^{iS/\hbar}$$

will be $O(\hbar)$ provided that S satisfies the Hamilton-Jacobi equation:

$$H\left(x_1, \dots, x_n, \frac{\partial S}{\partial x_1}, \dots, \frac{\partial S}{\partial x_n}\right) = \frac{\|\nabla S(x)\|^2}{2m} + V(x) = E.$$

Since φ is of order zero in \hbar , while $(\hat{H} - E)\varphi = O(\hbar)$, the ansatz φ again constitutes a first-order approximate solution to the time-independent Schrödinger equation.

We will call a phase function $S: \mathbb{R}^n \rightarrow \mathbb{R}$ **admissible** if it satisfies the Hamilton-Jacobi equation. The image $L = \text{im}(dS)$ of the differential of an admissible phase function S is characterized by three geometric properties:

1. L is an n -dimensional submanifold of $H^{-1}(E)$.
2. The pull-back to L of the form $\alpha_n = \sum_j p_j dq_j$ on \mathbb{R}^{2n} is exact.
3. The restriction of the canonical projection $\pi: T^*\mathbb{R}^n \rightarrow \mathbb{R}^n$ to L induces a diffeomorphism $L \simeq \mathbb{R}^n$. In other words, L is *projectable*.

While many of the basic constructions of microlocal analysis are motivated by operations on these projectable submanifolds of $T^*\mathbb{R}^n \simeq \mathbb{R}^{2n}$, applications of the theory require us to extend the constructions to more general n -dimensional submanifolds of \mathbb{R}^{2n} satisfying only a weakened version of condition (2) above, in which “exact” is replaced by “closed”. Such submanifolds are called **lagrangian**.

For example, the level sets for the 1-dimensional harmonic oscillator are lagrangian submanifolds in the phase plane. A regular level curve of the hamiltonian is an ellipse L . Since L is 1-dimensional, the pull-back to L of the form $p dq$ is closed, but the integral of $p dq$ around the curve equals the enclosed nonzero area, so its pull-back to L is not exact. It is also clear that the curve fails to project diffeomorphically onto \mathbb{R} . From the classical standpoint, the behavior of an oscillator is nevertheless completely described by its trajectory, suggesting that in general the state of a system should be represented by the submanifold L (projectable or not) rather than by the phase function S . This idea, which we will clarify later, is the starting point of the geometrical approach to microlocal analysis.

For now, we want to note an important relationship between lagrangian submanifolds of \mathbb{R}^{2n} and hamiltonian flows. Recall that to a function $H : \mathbb{R}^{2n} \rightarrow \mathbb{R}$, Hamilton's equations associate the vector field

$$X_H = \dot{q} \frac{\partial}{\partial q} + \dot{p} \frac{\partial}{\partial p} = \sum_j \frac{\partial H}{\partial p_j} \frac{\partial}{\partial q_j} - \frac{\partial H}{\partial q_j} \frac{\partial}{\partial p_j}.$$

A simple computation shows that X_H and the form α_n are related by the equation

$$X_H \lrcorner d\alpha_n = -dH,$$

i.e.

$$d\alpha_n(X_H, v) = -dH(v)$$

for every tangent vector v . If L is a lagrangian submanifold of a level set of H , then TL lies in the kernel of dH at all points of L , or, in other words, the 2-form $d\alpha_n$ vanishes on the subspace of $T_p\mathbb{R}^{2n}$ generated by T_pL and $X_H(p)$ for each $p \in L$. The restriction of $d\alpha_n$ to the tangent space $T_p\mathbb{R}^{2n}$ of \mathbb{R}^{2n} at any point p defines a nondegenerate, skew-symmetric bilinear form, and thus, as we will see in the next chapter, subspaces of $T_p\mathbb{R}^{2n}$ on which $d\alpha_n$ vanishes can be at most n -dimensional. These remarks imply that X_H is tangent to L , and we have the following result.

Hamilton-Jacobi theorem . *A function $H : \mathbb{R}^{2n} \rightarrow \mathbb{R}$ is locally constant on a lagrangian submanifold $L \subset \mathbb{R}^{2n}$ if and only if the hamiltonian vector field X_H is tangent to L .*

If the lagrangian submanifold L is locally closed, this theorem implies that L is invariant under the flow of X_H .

2.2 The WKB approximation

Returning to our WKB ansatz for a stationary-state solution of the Schrödinger equation, we recall that if $S : \mathbb{R}^n \rightarrow \mathbb{R}$ is an admissible phase function, then $\varphi(x) = e^{iS(x)/\hbar}$ satisfies

$$(\hat{H} - E)\varphi = O(\hbar).$$

Up to terms of order \hbar , in other words, φ is an eigenfunction of \hat{H} with eigenvalue E .

There is no way to improve the order of approximation simply by making a better choice of S . It is also clear on physical grounds that our ansatz for φ is too restrictive because it satisfies $|\varphi(x)| = 1$ for all x . In quantum mechanics, the quantity $|\varphi(x)|^2$ represents the probability of the particle being at the position x , and there is no reason for this to be constant; in fact, it is at least intuitively plausible that a particle is more likely to be found where it moves more slowly, i.e., where its potential energy is higher. We may therefore hope to find a better approximate solution by multiplying φ by an “amplitude function” a

$$\varphi(x) = e^{iS(x)/\hbar} a(x).$$

If S is again an admissible phase function, we now obtain:

$$(\hat{H} - E)\varphi = -\frac{1}{2m} \left[i\hbar \left(a\Delta S + 2 \sum_j \frac{\partial a}{\partial x_j} \frac{\partial S}{\partial x_j} \right) + \hbar^2 \Delta a \right] e^{iS/\hbar}.$$

If a is chosen to kill the coefficient of \hbar on the right, then φ will be an eigenfunction of \hat{H} modulo terms of order $O(\hbar^2)$. This condition on a is known as the **homogeneous transport equation**:

$$a\Delta S + 2 \sum_j \frac{\partial a}{\partial x_j} \frac{\partial S}{\partial x_j} = 0.$$

If S is an admissible phase function, and a is an amplitude which satisfies the homogeneous transport equation, then the second-order solution $\varphi = e^{iS/\hbar}a$ is called the **semi-classical approximation**.

Example 2.1 In 1 dimension, the homogeneous transport equation amounts to

$$aS'' + 2a'S' = 0.$$

Solving this equation directly, we obtain

$$\begin{aligned} a^2S'' + 2aa'S' &= (a^2S')' = 0 \\ \Rightarrow a &= \frac{c}{\sqrt{S'}} \end{aligned}$$

for some constant c . Since S is assumed to satisfy the Hamilton-Jacobi equation, we have $S' = \sqrt{2m(E - V)}$, and thus

$$a = \frac{c}{[4(E - V)]^{\frac{1}{4}}}.$$

If $E > V(x)$ for all $x \in \mathbb{R}$, then a is a smooth solution to the homogeneous transport equation. Notice that $a = |\varphi|$ is largest where V is largest, as our physical reasoning predicted.

Since the expression above for a does not depend explicitly on the phase function S , we might naively attempt to use the same formula when $\text{im}(dS)$ is replaced by a non-projectable lagrangian submanifold of $H^{-1}(E)$. Consider, for example, the unbounded potential $V(x) = x^2$ in the case of the harmonic oscillator. For $|x| < \sqrt{E}$, the function a is still well-defined up to a multiplicative constant. At $|x| = \sqrt{E}$, however, a has (asymptotic) singularities; observe that these occur precisely at the projected image of those points of L where the projection itself becomes singular. Outside the interval $|x| \leq \sqrt{E}$, the function a assumes complex values.

△

To generate better approximate solutions to the eigenfunction problem, we can extend the procedure above by adding to the original amplitude $a = a_0$ certain appropriately chosen functions of higher order in \hbar . Consider the next approximation

$$\varphi = e^{iS/\hbar}(a_0 + a_1\hbar).$$

Assuming that $e^{iS/\hbar}a_0$ is a semi-classical approximate solution, we obtain:

$$(\hat{H} - E)\varphi = -\frac{1}{2m} \left[i\hbar^2 \left(a_1 \Delta S + 2 \sum_j \frac{\partial a_1}{\partial x_j} \frac{\partial S}{\partial x_j} - i\Delta a_0 \right) + \hbar^3 \Delta a_1 \right] e^{iS(x)/\hbar}.$$

Evidently, φ will be a solution of the time-independent Schrödinger equation modulo terms of order $O(\hbar^3)$ provided that a_1 satisfies the **inhomogeneous transport equation**

$$a_1 \Delta S + 2 \sum_j \frac{\partial a_1}{\partial x_j} \frac{\partial S}{\partial x_j} = i\Delta a_0.$$

In general, a solution to the eigenfunction problem modulo terms of order $O(\hbar^n)$ is given by a WKB ansatz of the form

$$\varphi = e^{iS/\hbar}(a_0 + a_1\hbar + \dots + a_n\hbar^n),$$

where S satisfies the Hamilton-Jacobi equation, a_0 satisfies the homogeneous transport equation, and for each $k > 0$, the function a_k satisfies the inhomogeneous transport equation:

$$a_k \Delta S + 2 \sum_j \frac{\partial a_k}{\partial x_j} \frac{\partial S}{\partial x_j} = i\Delta a_{k-1}.$$

This situation can be described in the terminology of asymptotic series as follows. By an \hbar -dependent function f_\hbar on \mathbb{R}^n we will mean a function $f : \mathbb{R}^n \times \mathbb{R}_+ \rightarrow \mathbb{C}$, where \hbar is viewed as a parameter \hbar ranging in \mathbb{R}_+ . Such a function is said to be represented by a formal asymptotic expansion of the form $\sum_{k=0}^{\infty} a_k \hbar^k$, where each coefficient a_k is a smooth complex-valued function on \mathbb{R}^n , if, for each $K \in \mathbb{Z}_+$, the difference

$$f_\hbar - \sum_{k=0}^K a_k \hbar^k$$

is $O(\hbar^{K+1})$ locally uniformly in x . When f_\hbar admits such an expansion, its coefficients a_k are uniquely determined. It is obvious that any \hbar -dependent function which extends smoothly to $\hbar = 0$ is represented by an asymptotic series, and a theorem of E.Borel (see [28, p.28]) tells us that, conversely, any asymptotic series can be “summed” to yield such a function. The **principal part** of an asymptotic series $\sum_{k=0}^{\infty} a_k \hbar^k$ is defined as its first term which is not identically zero as a function of x . The **order** of a is the index of its principal part.

If we consider as equivalent any two \hbar -dependent functions whose difference is $O(\hbar^\infty)$, i.e. $O(\hbar^k)$ for all k , then each asymptotic series determines a unique equivalence class. A WKB “solution” to the eigenfunction problem $\hat{H}\varphi = E\varphi$ is then an equivalence class of functions of the form

$$\varphi = e^{iS/\hbar}a,$$

where S is an admissible phase function and a is an \hbar -dependent function represented by a formal asymptotic series

$$a \sim \sum_{k=p}^{\infty} a_k \hbar^k$$

with the property that its principal part a_p satisfies the homogeneous transport equation

$$a_p \Delta S + 2 \sum_j \frac{\partial a_p}{\partial x_j} \frac{\partial S}{\partial x_j} = 0,$$

and for $k > p$, the a_k satisfy the recursive transport equations:

$$a_k \Delta S + 2 \sum_j \frac{\partial a_k}{\partial x_j} \frac{\partial S}{\partial x_j} = i \Delta a_{k-1}.$$

This means that the \hbar -dependent function φ (or any function equivalent to it) satisfies the Schrödinger equation up to terms of order $O(\hbar^\infty)$.

Geometry of the transport equation

In Section 2.1, we saw that a first-order WKB approximate solution $\varphi = e^{iS/\hbar}$ to the time-independent Schrödinger equation depended on the choice of an admissible phase function, i.e., a function S satisfying the Hamilton-Jacobi equation $H(x, \frac{\partial S}{\partial x}) = E$. The generalized or geometric version of such a solution was a lagrangian submanifold of the level set $H^{-1}(E)$. We now wish to interpret and generalize in a similar way the semi-classical approximation with its amplitude included.

To begin, suppose that a is a function on \mathbb{R}^n which satisfies the homogeneous transport equation:

$$a \Delta S + 2 \sum_j \frac{\partial a}{\partial x_j} \frac{\partial S}{\partial x_j} = 0.$$

After multiplying both sides of this equation by a , we can rewrite it as:

$$\sum_j \frac{\partial}{\partial x_j} \left(a^2 \frac{\partial S}{\partial x_j} \right) = 0,$$

which means that the divergence of the vector field $a^2 \nabla S$ is zero. Rather than considering the transport equation as a condition on the vector field $a^2 \nabla S$ (on \mathbb{R}^n) *per se*, we can lift all of this activity to the lagrangian submanifold $L = \text{im}(dS)$. Notice first that the restriction to L of the hamiltonian vector field associated to $H(q, p) = \sum p_i^2/2 + V(q)$ is

$$X_H|_L = \sum_j \left(\frac{\partial S}{\partial x_j} \frac{\partial}{\partial q_j} - \frac{\partial V}{\partial q_j} \frac{\partial}{\partial p_j} \right).$$

The projection $X_H^{(x)}$ of $X_H|_L$ onto \mathbb{R}^n (the (x) reminds us of the coordinate x on \mathbb{R}^n) therefore coincides with ∇S , and so the homogeneous transport equation says that $a^2 X_H^{(x)}$ is divergence-free for the canonical density $|dx| = |dx_1 \wedge \cdots \wedge dx_n|$ on \mathbb{R}^n . But it is better to reformulate this condition as:

$$\mathcal{L}_{X_H^{(x)}}(a^2 |dx|) = 0;$$

that is, we transfer the factor of a^2 from the vector field $X_H^{(x)} = \nabla S$ to the density $|dx|$. Since X_H is tangent to L by the Hamilton-Jacobi theorem, and since the Lie derivative is invariant under diffeomorphism, this equation is satisfied if and only if the pull-back of $a^2|dx|$ to L via the projection π is invariant under the flow of X_H .

This observation, together with the fact that it is the *square* of a which appears in the density $\pi^*(a^2|dx|)$, suggests that a solution of the homogeneous transport equation should be represented geometrically by a *half-density* on L , invariant by X_H . (See Appendix A for a discussion of densities of fractional order.)

In other words, a (geometric) semi-classical state should be defined as a lagrangian submanifold L of \mathbb{R}^{2n} equipped with a half-density a . Such a state is stationary when L lies in a level set of the classical hamiltonian and a is invariant under its flow.

Example 2.2 Recall that in the case of the 1-dimensional harmonic oscillator, stationary classical states are simply those lagrangian submanifolds of \mathbb{R}^2 which coincide with the regular level sets of the classical hamiltonian $H(q, p) = (p^2 + kq^2)/2$. There is a unique (up to a constant) invariant volume element for the hamiltonian flow of H on each level curve H . Any such level curve L , together with a square root of this volume element, constitutes a semi-classical stationary state for the harmonic oscillator.

△

Notice that while a solution to the homogeneous transport equation in the case of the 1-dimensional harmonic oscillator was necessarily singular (see Example 2.1), the semi-classical state described in the preceding example is a perfectly smooth object everywhere on the lagrangian submanifold L . The singularities arise only when we try to transfer the half-density from L down to configuration space. Another substantial advantage of the geometric interpretation of the semi-classical approximation is that the concept of an invariant half-density depends only on the hamiltonian vector field X_H and not on the function S , so it makes sense on any lagrangian submanifold of \mathbb{R}^{2n} lying in a level set of H .

This discussion leads us to another change of viewpoint, namely that the quantum states themselves should be represented, not by functions, but by half-densities on configuration space \mathbb{R}^n , i.e. elements of the intrinsic Hilbert space $\mathfrak{H}_{\mathbb{R}^n}$ (see Appendix A). Stationary states are then eigenvectors of the Schrödinger operator \hat{H} , which is defined on the space of smooth half-densities in terms of the old Schrödinger operator on functions, which we will denote momentarily as \hat{H}_{fun} , by the equation

$$\hat{H}(a|dx|^{1/2}) = (\hat{H}_{\text{fun}}a)|dx|^{1/2}.$$

From this new point of view, we can express the result of our analysis as follows:

*If S is an admissible phase function and a is a half-density on $L = \text{im}(dS)$ which is invariant under the flow of the classical hamiltonian, then $e^{iS/\hbar}(dS)^*a$ is a second-order approximate solution to the time-independent Schrödinger equation.*

In summary, we have noted the following correspondences between classical and quantum mechanics:

Object	Classical version	Quantum version
basic space	\mathbb{R}^{2n}	$\mathfrak{H}_{\mathbb{R}^n}$
state	lagrangian submanifold of \mathbb{R}^{2n} with half-density	half-density on \mathbb{R}^n
time-evolution	Hamilton's equations	Schrödinger equation
generator of evolution	function H on \mathbb{R}^{2n}	operator \hat{H} on smooth half-densities
stationary state	lagrangian submanifold in level set of H with invariant half-density	eigenvector of \hat{H}

Proceeding further, we could attempt to interpret a solution of the recursive system of inhomogeneous transport equations on \mathbb{R}^n as an asymptotic half-density on L in order to arrive at a geometric picture of a complete WKB solution to the Schrödinger equation. This, however, involves some additional difficulties, notably the lack of a geometric interpretation of the inhomogeneous transport equations, which lie beyond the scope of these notes. Instead, we will focus on two aspects of the semi-classical approximation. First, we will extend the geometric picture presented above to systems with more general phase spaces. This will require the concept of symplectic manifold, which is introduced in the following chapter. Second, we will “quantize” semi-classical states in these symplectic manifolds. Specifically, we will try to construct a space of quantum states corresponding to a general classical phase space. Then we will try to construct asymptotic quantum states corresponding to half-densities on lagrangian submanifolds. In particular, we will start with an invariant half-density on a (possibly non-projectable) lagrangian submanifold of \mathbb{R}^{2n} and attempt to use this data to construct an explicit semi-classical approximate solution to Schrödinger's equation on \mathbb{R}^n .

3 Symplectic Manifolds

In this chapter, we will introduce the notion of a symplectic structure on a manifold, motivated for the most part by the situation in \mathbb{R}^{2n} . While some discussion will be devoted to certain general properties of symplectic manifolds, our main goal at this point is to develop the tools needed to extend the hamiltonian viewpoint to phase spaces associated to general finite-dimensional configuration spaces, i.e. to cotangent bundles. More general symplectic manifolds will reappear as the focus of more sophisticated quantization programs in later chapters. We refer to [6, 29, 63] for thorough discussions of the topics in this chapter.

3.1 Symplectic structures

In Section 2.1, a *lagrangian submanifold* of \mathbb{R}^{2n} was defined as an n -dimensional submanifold $L \subset \mathbb{R}^{2n}$ on which the exterior derivative of the form $\alpha_n = \sum p_i dq_i$ vanishes; to a function $H: \mathbb{R}^{2n} \rightarrow \mathbb{R}$, we saw that Hamilton's equations associate a vector field X_H on \mathbb{R}^{2n} satisfying

$$X_H \lrcorner d\alpha_n = -dH.$$

Finally, our proof of the Hamilton-Jacobi theorem relied on the *nondegeneracy* of the 2-form $d\alpha_n$. These points already indicate the central role played by the form $d\alpha_n$ in the study of hamiltonian systems in \mathbb{R}^{2n} ; the correct generalization of the hamiltonian picture to arbitrary configuration spaces relies similarly on the use of 2-forms with certain additional properties. In this section, we first study such forms pointwise, collecting pertinent facts about nondegenerate, skew-symmetric bilinear forms. We then turn to the definition of symplectic manifolds.

Linear symplectic structures

Suppose that V is a real, m -dimensional vector space. A bilinear form $\omega: V \times V \rightarrow \mathbb{R}$ gives rise to a linear map

$$\tilde{\omega}: V \rightarrow V^*$$

defined by contraction:

$$\tilde{\omega}(x)(y) = \omega(x, y).$$

The ω -orthogonal to a subspace $W \subset V$ is defined as

$$W^\perp = \{x \in V : W \subset \ker \tilde{\omega}(x)\}.$$

If $\tilde{\omega}$ is an isomorphism, or in other words if $V^\perp = \{0\}$, then the form ω is said to be *nondegenerate*; if in addition ω is skew-symmetric, then ω is called a *linear symplectic structure* on V . A linear endomorphism of a symplectic vector space (V, ω) which preserves the form ω is called a **linear symplectic transformation**, and the group of all such transformations is denoted by $Sp(V)$.

Example 3.1 If E is any real n -dimensional vector space with dual E^* , then a linear symplectic structure on $V = E \oplus E^*$ is given by

$$\omega((x, \lambda), (x', \lambda')) = \lambda'(x) - \lambda(x').$$

With respect to a basis $\{x_i\}$ of E and a dual basis $\{\lambda_i\}$ of E^* , the form ω is represented by the matrix

$$\omega = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}.$$

It follows that if a linear operator on V is given by the real $2n \times 2n$ matrix

$$T = \begin{pmatrix} A & B \\ C & D \end{pmatrix},$$

then T is symplectic provided that $A^t C, B^t D$ are symmetric, and $A^t D - C^t B = I$. Note in particular that these conditions are satisfied if $A \in GL(E)$, $D = (A^t)^{-1}$, and $B = C = 0$, and so $GL(E)$ is isomorphic to a subgroup $Gl(E)$ of $Sp(V)$. More generally, if $K: E \rightarrow F$ is an isomorphism, then the association

$$(x, \lambda) \mapsto (Kx, (K^{-1})^* \lambda)$$

defines a linear symplectomorphism between $E \oplus E^*$ and $F \oplus F^*$ equipped with these linear symplectic structures.

△

Since the determinant of a skew-symmetric $m \times m$ matrix is zero if m is odd, the existence of a linear symplectic structure on a vector space V implies that V is necessarily even-dimensional and therefore admits a complex structure, i.e. a linear endomorphism J such that $J^2 = -I$. A complex structure is said to be *compatible* with a symplectic structure on V if

$$\omega(Jx, Jy) = \omega(x, y)$$

and

$$\omega(x, Jx) > 0$$

for all $x, y \in V$. In other words, J is compatible with ω (we also call it ω -compatible) if $J: V \rightarrow V$ is a linear symplectomorphism and $g_J(\cdot, \cdot) = \omega(\cdot, J\cdot)$ defines a symmetric, positive-definite bilinear form on V .

Theorem 3.2 *Every symplectic vector space (V, ω) admits a compatible complex structure.*

Proof. Let $\langle \cdot, \cdot \rangle$ be a symmetric, positive-definite inner product on V , so that ω is represented by an invertible skew-adjoint operator $K: V \rightarrow V$; i.e.

$$\omega(x, y) = \langle Kx, y \rangle.$$

The operator K admits a polar decomposition $K = RJ$, where $R = \sqrt{KK^t}$ is positive-definite symmetric, $J = R^{-1}K$ is orthogonal, and $RJ = JR$. From the skew-symmetry of K it follows that $J^t = -J$, and so $J^2 = -JJ^t = -\text{id}$; i.e., J is a complex structure on V .

To see that J is ω -compatible, first note that

$$\omega(Jx, Jy) = \langle KJx, Jy \rangle = \langle JKx, Jy \rangle = \langle Kx, y \rangle = \omega(x, y).$$

Also,

$$\omega(x, Jx) = \langle Kx, Jx \rangle = \langle JRx, Jx \rangle = \langle Rx, x \rangle > 0,$$

since R and $\langle \cdot, \cdot \rangle$ are positive-definite. □

Corollary 3.3 *The collection \mathcal{J} of ω -compatible complex structures on a symplectic vector space (V, ω) is contractible.*

Proof. The association $J \mapsto g_J$ described above defines a continuous map from \mathcal{J} into the space \mathcal{P} of symmetric, positive-definite bilinear forms on V . By the uniqueness of the polar decomposition, it follows that the map which assigns to a form $\langle \cdot, \cdot \rangle$ the complex structure J constructed in the preceding proof is continuous, and the composition $\mathcal{J} \rightarrow \mathcal{P} \rightarrow \mathcal{J}$ of these maps equals the identity on \mathcal{J} . Since \mathcal{P} is contractible, this implies the corollary. □

If J is ω -compatible, a hermitian structure on V is defined by

$$\langle \cdot, \cdot \rangle = g_J(\cdot, \cdot) + i\omega(\cdot, \cdot).$$

As is easily checked, a linear transformation $T \in GL(V)$ which preserves any two of the structures ω, g_J, J on V preserves the third and therefore preserves the hermitian structure. In terms of the automorphism groups $Sp(V), GL(V, J), O(V)$, and $U(V)$ of ω, J, g , and $\langle \cdot, \cdot \rangle$, this means that the intersection of any two of $Sp(V), GL(V, J), O(V)$ equals $U(V)$.

To determine the Lie algebra $\mathfrak{sp}(\mathfrak{V})$ of $Sp(V)$, consider a 1-parameter family of maps e^{tA} associated to some linear map $A: V \rightarrow V$. For any $v, w \in V$, we have

$$\left. \frac{d}{dt} \right|_{t=0} \omega(e^{tA}v, e^{tA}w) = \omega(Av, w) + \omega(v, Aw),$$

and so $A \in \mathfrak{sp}(\mathfrak{V})$ if and only if the linear map $\tilde{\omega} \circ A: V \rightarrow V^*$ is self-adjoint. Consequently, $\dim(V) = 2k$ implies $\dim(\mathfrak{sp}(\mathfrak{V})) = \dim(\mathfrak{Sp}(\mathfrak{V})) = \mathfrak{k}(2\mathfrak{k} + \mathfrak{1})$.

Distinguished subspaces

The ω -orthogonal to a subspace W of a symplectic vector space (V, ω) is called the *symplectic orthogonal* to W . From the nondegeneracy of the symplectic form, it follows that

$$W^{\perp\perp} = W \quad \text{and} \quad \dim W^{\perp} = \dim V - \dim W$$

for any subspace $W \subset V$. Also,

$$(A + B)^\perp = A^\perp \cap B^\perp \quad \text{and} \quad (A \cap B)^\perp = A^\perp + B^\perp$$

for any pair of subspaces A, B of V . In particular, $B^\perp \subset A^\perp$ whenever $A \subset B$.

Note that the symplectic orthogonal W^\perp might not be an algebraic complement to W . For instance, if $\dim W = 1$, the skew-symmetry of ω implies that $W \subset W^\perp$. More generally, any subspace contained in its orthogonal will be called *isotropic*. Dually, we note that if $\text{codim } W = 1$, then W^\perp is 1-dimensional, hence isotropic, and $W^\perp \subset W^{\perp\perp} = W$. In general, spaces W satisfying the condition $W^\perp \subset W$ are called *coisotropic* or *involutive*. Finally, if W is self-orthogonal, i.e. $W^\perp = W$, then the dimension relation above implies that $\dim W = \frac{1}{2} \dim V$. Any self-orthogonal subspace is simultaneously isotropic and coisotropic, and is called *lagrangian*.

According to these definitions, a subspace $W \subset V$ is isotropic if the restriction of the symplectic form to W is identically zero. At the other extreme, the restriction of ω to certain subspaces $Z \subset V$ may again be nondegenerate; this is equivalent to saying that $Z \cap Z^\perp = \{0\}$ or $Z + Z^\perp = V$. Such subspaces are called *symplectic*.

Example 3.4 In $E \oplus E^*$ with its usual symplectic structure, both E and E^* are lagrangian subspaces. It also follows from the definition of this structure that the graph of a linear map $B: E \rightarrow E^*$ is a lagrangian subspace of $E \oplus E^*$ if and only if B is self-adjoint.

If (V, ω) is a symplectic vector space, we denote by $V \oplus \bar{V}$ the vector space $V \oplus V$ equipped with the symplectic structure $\omega \oplus -\omega$. If $T: V \rightarrow V$ is a linear symplectic map, then the graph of T is a lagrangian subspace of $V \oplus \bar{V}$.

The kernel of a nonzero covector $\alpha \in V^*$ is a codimension-1 coisotropic subspace $\ker \alpha$ of V whose symplectic orthogonal $(\ker \alpha)^\perp$ is the distinguished 1-dimensional subspace of $\ker \alpha$ spanned by $\tilde{\omega}^{-1}(\alpha)$.

△

Example 3.5 Suppose that (V, ω) is a $2n$ -dimensional symplectic vector space and $W \subset V$ is any isotropic subspace with $\dim(W) < n$. Since $2n = \dim(W) + \dim(W^\perp)$, there exists a nonzero vector $w \in W^\perp \setminus W$. The subspace W' of V spanned by $W \cup \{w\}$ is then isotropic and $\dim(W') = \dim(W) + 1$. From this observation it follows that for every isotropic subspace W of a (finite-dimensional) symplectic vector space V which is not lagrangian, there exists an isotropic subspace W' of V which properly contains W . Beginning with any 1-dimensional subspace of V , we can apply this remark inductively to conclude that every finite-dimensional symplectic vector space contains a lagrangian subspace.

△

Various subspaces of a symplectic vector space are related as follows.

Lemma 3.6 *If L is a lagrangian subspace of a symplectic vector space V , and $A \subset V$ is an arbitrary subspace, then:*

1. $L \subset A$ if and only if $A^\perp \subset L$.
2. L is transverse to A if and only if $L \cap A^\perp = \{0\}$.

Proof. Statement (1) follows from the properties of the operation $^\perp$ and the equation $L = L^\perp$. Similarly, $L + A = V$ if and only if $(L + A)^\perp = L \cap A^\perp = \{0\}$, proving statement (2).

□

Example 3.7 Note that statement (1) of Lemma 3.6 implies that if $L \subset A$, then $A^\perp \subset A$, and so A is a coisotropic subspace. Conversely, if A is coisotropic, then A^\perp is isotropic, and Example 3.5 implies that there is a lagrangian subspace L with $A^\perp \subset L$. Passing to orthogonals, we have $L \subset A$. Thus, a subspace $C \subset V$ is coisotropic if and only if it contains a lagrangian subspace.

Suppose that V is a symplectic vector space with an isotropic subspace I and a lagrangian subspace L such that $I \cap L = 0$. If $W \subset L$ is any complementary subspace to $I^\perp \cap L$, then $I + L \subset W + I^\perp$, and so $W^\perp \cap I \subset I^\perp \cap L$. Thus, $W^\perp \cap I \subset I \cap L = 0$. Since $I^\perp \cap W = 0$ by our choice of W , it follows that $I + W$ is a symplectic subspace of V .

△

A pair L, L' of transverse lagrangian subspaces of V is said to define a **lagrangian splitting** of V . In this case, the map $\tilde{\omega}$ defines an isomorphism $L' \simeq L^*$, which in turn gives rise to a linear symplectomorphism between V and $L \oplus L^*$ equipped with its canonical symplectic structure (see Example 3.1). If J is a ω -compatible complex structure on V and $L \subset V$ a lagrangian subspace, then L, JL is a lagrangian splitting. By Example 3.5, every symplectic vector space contains a lagrangian subspace, and since every n -dimensional vector space is isomorphic to \mathbb{R}^n , the preceding remarks prove the following linear “normal form” result:

Theorem 3.8 *Every $2n$ -dimensional symplectic vector space is linearly symplectomorphic to $(\mathbb{R}^{2n}, \omega_n)$.*

Theorem 3.2 also implies the following useful result.

Lemma 3.9 *Suppose that V is a symplectic vector space with a ω -compatible complex structure J and let $T_\varepsilon: V \rightarrow V$ be given by $T_\varepsilon(x) = x + \varepsilon Jx$.*

1. *If L, L' are any lagrangian subspaces of V , then $L_\varepsilon = T_\varepsilon(L)$ is a lagrangian subspace transverse to L' for small $\varepsilon > 0$.*
2. *For any two lagrangian subspaces L, L' of V , there is a lagrangian subspace L'' transverse to both L' and L .*

Proof. It is easy to check that T_ε is a *conformal* linear symplectic map, i.e. an isomorphism of V satisfying $\omega(T_\varepsilon x, T_\varepsilon y) = (1 + \varepsilon^2) \omega(x, y)$. Thus, L_ε is a lagrangian subspace for all $\varepsilon > 0$. Using the inner-product g_J on V induced by J , we can choose orthonormal bases

$\{v_i\}, \{w_i\}$ of L' and L , respectively, so that for $i = 1, \dots, k$, the vectors $v_i = w_i$ span $L \cap L'$. Then $\{w_i + \varepsilon Jw_i\}$ form a basis of L_ε , and L', L_ε are transverse precisely when the matrix $M = \{\omega(v_i, w_j + \varepsilon Jw_j)\} = \{\omega(v_i, w_j) + \varepsilon g_J(v_i, w_j)\}$ is nonsingular. Our choice of bases implies that

$$M = \begin{pmatrix} \varepsilon \cdot \text{id} & 0 \\ 0 & A + \varepsilon \cdot B \end{pmatrix}$$

where $A = \{\omega(v_i, w_j)\}_{i,j=k+1}^n$ and B is some $(n-k) \times (n-k)$ matrix. Setting $I = \text{span}\{v_i\}_{i=k+1}^n$ and $W = \text{span}\{w_i\}_{i=k+1}^n$, we can apply Example 3.7 to conclude that A is nonsingular, and assertion (1) follows.

To prove (2), observe that for small $\varepsilon > 0$, the lagrangian subspace L_ε is transverse to L, L' by (1).

□

In fact, the statement of preceding lemma can be improved as follows. Let $\{L_i\}$ be a countable family of lagrangian subspaces, and let A_i, B_i be the matrices obtained with respect to L as in the proof above. For each i , the function $t \mapsto \det(A_i + tB_i)$ is a nonzero polynomial and therefore has finitely many zeros. Consequently, the lagrangian subspace $T_t(L)$ is transverse to all L_i for almost every $t \in \mathbb{R}$.

The lagrangian grassmannian

The collection of all unoriented lagrangian subspaces of a $2n$ -dimensional symplectic vector space V is called the lagrangian grassmannian $\mathcal{L}(V)$ of V . A natural action of the group $Sp(V)$ on $\mathcal{L}(V)$, denoted $j: Sp(V) \times \mathcal{L}(V) \rightarrow \mathcal{L}(V)$ is defined by $j(T, L) = j_L(T) = T(L)$.

Lemma 3.10 *The unitary group associated to an ω -compatible complex structure J on V acts transitively on $\mathcal{L}(V)$.*

Proof. For arbitrary $L_1, L_2 \in \mathcal{L}(V)$, an orthogonal transformation $L_1 \rightarrow L_2$ induces a symplectic transformation $L_1 \oplus L_1^* \rightarrow L_2 \oplus L_2^*$ in the manner of Example 3.1, which in turn gives rise to a unitary transformation $L_1 \oplus JL_1 \rightarrow L_2 \oplus JL_2$ mapping L_1 onto L_2 .

□

The stabilizer of $L \in \mathcal{L}(V)$ under the $U(V)$ -action is evidently the orthogonal subgroup of $Gl(L)$ defined with respect to the inner-product and splitting $L \oplus JL$ of V induced by J (see Example 3.1). Thus, a (non-canonical) identification of the lagrangian grassmannian with the homogeneous space $U(n)/O(n)$ is obtained from the map

$$U(V) \xrightarrow{j_L} \mathcal{L}(V).$$

The choice of J also defines a complex determinant $U(V) \xrightarrow{\det^2 J} S^1$, which induces a fibration $\mathcal{L}(V) \rightarrow S^1$ with 1-connected fiber $SU(n)/SO(n)$, giving an isomorphism of fundamental groups

$$\pi_1(\mathcal{L}(V)) \simeq \pi_1(S^1) \simeq \mathbb{Z}.$$

This isomorphism does not depend on the choices of J and L made above. Independence of J follows from the fact that \mathcal{J} is connected (Corollary 3.3). On the other hand, connectedness of the unitary group together with Lemma 3.10 gives independence of L .

Passing to homology and dualizing, we obtain a *natural* homomorphism

$$H^1(S^1; \mathbb{Z}) \rightarrow H^1(\mathcal{L}(V); \mathbb{Z}).$$

The image of the canonical generator of $H^1(S^1; \mathbb{Z})$ under this map is called the *universal Maslov class*, μ_V . The result of the following example will be useful when we extend our discussion of the Maslov class from vector spaces to vector bundles.

Example 3.11 If (V, ω) is any symplectic vector space with ω -compatible complex structure J and lagrangian subspace L , then a check of the preceding definitions shows that

$$m_L(j(T, L')) = m_L(L') \cdot \det_J^2(T)$$

for any $T \in U(V)$ and $L' \in \mathcal{L}(V)$. (Recall that $j: Sp(V) \times \mathcal{L}(V) \rightarrow \mathcal{L}(V)$ denotes the natural action of $Sp(V)$ on $\mathcal{L}(V)$).

Now consider any topological space M . If $f_1, f_2: M \rightarrow \mathcal{L}(V)$ are continuous maps, then the definition of the universal Maslov class shows that $(f_1^* - f_2^*)\mu_V$ equals the pull-back of the canonical generator of $H^1(S^1; \mathbb{R})$ by the map

$$(m_L \circ f_1)(m_L \circ f_2)^{-1}.$$

(Here we use the fact that when S^1 is identified with the unit complex numbers, the multiplication map $S^1 \times S^1 \rightarrow S^1$ induces the diagonal map $H^1(S^1) \rightarrow H^1(S^1) \oplus H^1(S^1) \simeq H^1(S^1 \times S^1)$ on cohomology). If $T: M \rightarrow Sp(V)$ is any map, we set $(T \cdot f_i) = j(T, f_i)$. Since $Sp(V)$ deformation retracts onto $U(V)$, it follows that T is homotopic to a map $T': M \rightarrow U(V)$, and so $((T \cdot f_1)^* - (T \cdot f_2)^*)\mu_V$ is obtained via pull-back by

$$(m_L \circ (T' \cdot f_1))(m_L \circ (T' \cdot f_2))^{-1}.$$

From the first paragraph, it follows that this product equals $(m_L \circ f_1)(m_L \circ f_2)^{-1}$, from which we conclude that

$$(f_1^* - f_2^*)\mu_V = ((T \cdot f_1)^* - (T \cdot f_2)^*)\mu_V.$$

△

Symplectic manifolds

To motivate the definition of a symplectic manifold, we first recall some features of the differential form $-d\alpha_n = \omega_n = \sum_{j=1}^n dq_j \wedge dp_j$ which appeared in our earlier discussion. First, we note that $\sum_{j=1}^n dq_j \wedge dp_j$ defines a linear symplectic structure on the tangent space of \mathbb{R}^{2n} at each point. In fact:

$$\omega_n \left(\frac{\partial}{\partial q_j}, \frac{\partial}{\partial p_k} \right) = \delta_{jk} \quad \omega_n \left(\frac{\partial}{\partial q_j}, \frac{\partial}{\partial q_k} \right) = 0 \quad \omega_n \left(\frac{\partial}{\partial p_j}, \frac{\partial}{\partial p_k} \right) = 0$$

and so

$$\tilde{\omega}_n \left(\frac{\partial}{\partial q_j} \right) = dp_j \quad \tilde{\omega}_n \left(\frac{\partial}{\partial p_j} \right) = -dq_j,$$

from which it is clear that $\tilde{\omega}_n$ is invertible.

Next, we recall that the hamiltonian vector field associated via Hamilton's equations to $H : \mathbb{R}^{2n} \rightarrow \mathbb{R}$ satisfies

$$X_H \lrcorner \omega_n = dH,$$

or in other words,

$$X_H = \tilde{\omega}_n^{-1}(dH),$$

so we see that the symplectic form ω_n is all that we need to obtain X_H from H . This description of the hamiltonian vector field leads immediately to the following two invariance results. First note that by the skew-symmetry of ω_n ,

$$\mathcal{L}_{X_H} H = X_H \cdot H = \omega_n(X_H, X_H) = 0,$$

implying that X_H is tangent to the level sets of H . This again reflects the fact that the flow of X_H preserves energy. Since ω_n is closed, we also have by Cartan's formula (see [1])

$$\mathcal{L}_{X_H} \omega_n = d(X_H \lrcorner \omega_n) + X_H \lrcorner d\omega_n = d^2 H = 0.$$

This equation implies that the flow of X_H preserves the form ω_n and therefore generalizes our earlier remark that the hamiltonian vector field associated to the 1-dimensional harmonic oscillator is divergence-free.

We now see what is needed to do hamiltonian mechanics on manifolds. A 2-form ω on a manifold P is a smooth family of bilinear forms on the tangent spaces of P . By assuming that each of these bilinear forms is nondegenerate, we guarantee that the equation $X_H = \tilde{\omega}^{-1}(dH)$ defines a hamiltonian vector field uniquely for any H . Computing the Lie derivative of H with respect to X_H

$$\mathcal{L}_{X_H} H = X_H \cdot H = \tilde{\omega}(X_H)(X_H) = \omega(X_H, X_H) = 0,$$

we see that the conservation of energy follows from the skew-symmetry of the form ω .

Finally, invariance of ω under the hamiltonian flow is satisfied if

$$\mathcal{L}_{X_H} \omega = d(X_H \lrcorner \omega) + X_H \lrcorner d\omega = 0.$$

Here, the term $d(X_H \lrcorner \omega) = d^2 H$ is automatically zero; to guarantee the vanishing of the second term, we impose the condition that ω be closed.

Thus we make the following definition:

Definition 3.12 *A symplectic structure on a manifold P is a closed, nondegenerate 2-form ω on P .*

The condition that ω be nondegenerate means that $\tilde{\omega}$ defines an isomorphism of vector bundles $TP \rightarrow T^*P$, or equivalently, that the top exterior power of ω is a volume form on P , or finally, that ω defines a linear symplectic structure on each tangent space of P .

An immediate example of a symplectic manifold is furnished by \mathbb{R}^{2n} with its standard structure $\omega_n = \sum_{j=1}^n dq_j \wedge dp_j$ (a differential form with constant coefficients and not just a single bilinear form). Darboux's theorem (Section 4.3) will tell us that this is the local model for the general case. In the next section, we will see that the cotangent bundle of any smooth manifold carries a natural symplectic structure.

Generalizing our earlier discussion of distinguished subspaces of a symplectic vector space, we call a submanifold $C \subset P$ *(co-)isotropic* provided that each tangent space $T_p C$ of C is a (co-)isotropic subspace of $T_p P$. When C is coisotropic, the subspaces $\mathcal{C}_p = (T_p C)^\perp$ comprise a subbundle $(TC)^\perp$ of TC known as the **characteristic distribution** of C . It is integrable because ω is closed. Of particular interest in our discussion will be *lagrangian* submanifolds of P , which are (co-)isotropic submanifolds of dimension $\frac{1}{2} \dim(P)$. More generally, if L is a smooth manifold of dimension $\frac{1}{2} \dim(P)$ and $\iota: L \rightarrow P$ is an immersion such that $\iota^* \omega = 0$, we will call the pair (L, ι) a **lagrangian immersion**.

Example 3.13 If $C \subset P$ is a hypersurface, then C is a coisotropic submanifold. A simple check of definitions shows that if $H: P \rightarrow \mathbb{R}$ is a smooth function having C as a regular level set, then the hamiltonian vector field X_H is tangent to the characteristic foliation of C .

If (L, ι) is a lagrangian immersion whose image is contained in C , then Lemma 3.6 implies that for each $p \in L$, the characteristic subspace $\mathcal{C}_{\iota(p)} \subset T_{\iota(p)} C$ is contained in $\iota_* T_p L$, and thus X_H induces a smooth, nonsingular vector field $X_{H,\iota}$ on L . In view of the remarks above, this assertion generalizes the Hamilton-Jacobi theorem (see the end of Section 2.1) to arbitrary symplectic manifolds and lagrangian immersions.

△

New symplectic manifolds can be manufactured from known examples by dualizing and by taking products. The **symplectic dual** of a manifold (P, ω) consists of the same underlying manifold endowed with the symplectic structure $-\omega$. Evidently P and its dual \bar{P} share the same (co-)isotropic submanifolds. Given two symplectic manifolds (P_1, ω_1) and (P_2, ω_2) , their product $P_1 \times P_2$ admits a symplectic structure given by the sum $\omega_1 \oplus \omega_2$. More explicitly, this form is the sum of the pull-backs of ω_1 and ω_2 by the projections of $P_1 \times P_2$ to P_1 and P_2 . As is easily verified, the product of (co-)isotropic submanifolds of P_1 and P_2 is a (co-)isotropic submanifold of $P_1 \times P_2$.

A **symplectomorphism** from (P_1, ω_1) to (P_2, ω_2) is a smooth diffeomorphism $f: P_1 \rightarrow P_2$ compatible with the symplectic structures: $f^* \omega_2 = \omega_1$. A useful connection among duals, products, and symplectomorphisms is provided by the following lemma.

Lemma 3.14 *A diffeomorphism $f: P_1 \rightarrow P_2$ between symplectic manifolds is a symplectomorphism if and only if its graph is a lagrangian submanifold of the product $P_2 \times \bar{P}_1$.*

The collection $\text{Aut}(P, \omega)$ of symplectomorphisms of P becomes an infinite-dimensional Lie group when endowed with the C^∞ topology (see [49]). In this case, the corresponding Lie

algebra is the space $\chi(P, \omega)$ of smooth vector fields X on P satisfying

$$\mathcal{L}_X \omega = 0.$$

Since $\mathcal{L}_X \omega = d(X \lrcorner \omega)$, the association $X \mapsto X \lrcorner \omega$ defines an isomorphism between $\chi(P, \omega)$ and the space of closed 1-forms on P ; those X which map to exact 1-forms are simply the hamiltonian vector fields on P . The elements of $\chi(P, \omega)$ are called **locally hamiltonian vector fields** or **symplectic vector fields**.

Example 3.15 A linear symplectic form ω on a vector space V induces a symplectic structure (also denoted ω) on V via the canonical identification of TV with $V \times V$. The symplectic group $Sp(V)$ then embeds naturally in $\text{Aut}(V, \omega)$, and the Lie algebra $\mathfrak{sp}(\mathfrak{V})$ identifies with the subalgebra of $\chi(V, \omega)$ consisting of vector fields of the form

$$X(v) = Av$$

for some $A \in \mathfrak{sp}(\mathfrak{V})$. Note that these are precisely the hamiltonian vector fields of the homogeneous quadratic polynomials on V , i.e. functions satisfying $Q(tv) = t^2Q(v)$ for all real t . Consequently, $\mathfrak{sp}(\mathfrak{V})$ is canonically identified with the space of such functions via the correspondence

$$A \leftrightarrow Q_A(v) = \frac{1}{2} \omega(Av, v).$$

△

Symplectic vector bundles

Since a symplectic form on a $2n$ -manifold P defines a smooth family of linear symplectic forms on the fibers of TP , the frame bundle of P can be reduced to a principal $Sp(n)$ bundle over P . More generally, any vector bundle $E \rightarrow B$ with this structure is called a **symplectic vector bundle**. Two symplectic vector bundles E, F are said to be symplectomorphic if there exists a vector bundle isomorphism $E \rightarrow F$ which preserves their symplectic structures.

Example 3.16 If $F \rightarrow B$ is any vector bundle, then the sum $F \oplus F^*$ carries a natural symplectic vector bundle structure, defined in analogy with Example 3.1.

△

With the aid of an arbitrary riemannian metric, the proof of Theorem 3.2 can be generalized by a fiberwise construction as follows.

Theorem 3.17 *Every symplectic vector bundle admits a compatible complex vector bundle structure.*

Example 3.18 Despite Theorem 3.17, there exist examples of symplectic manifolds which are not complex (the almost complex structure coming from the theorem cannot be made integrable), and of complex manifolds which are not symplectic. (See [27] and the numerous earlier references cited therein.) Note, however, that the Kähler form of any Kähler manifold is a symplectic form.

△

A **lagrangian subbundle** of a symplectic vector bundle E is a subbundle $L \subset E$ such that L_x is a lagrangian subspace of E_x for all $x \in B$. If E admits a lagrangian subbundle, then E is symplectomorphic to $L \oplus L^*$, and the frame bundle of E admits a further reduction to a principal $GL(n)$ bundle over B (compare Example 3.1).

Example 3.19 If L is a lagrangian submanifold of a symplectic manifold P , then the restricted tangent bundle $T_L P$ is a symplectic vector bundle over L , and $TL \subset T_L P$ is a lagrangian subbundle. Also note that if $C \subset P$ is any submanifold such that TC contains a lagrangian subbundle of $T_C P$, then C is coisotropic (see Lemma 3.6).

△

In general, the automorphism group of a symplectic vector bundle E does not act transitively on the lagrangian subbundles of E . Nevertheless, a pair of transverse lagrangian subbundles can be related as follows.

Theorem 3.20 *Let $E \rightarrow B$ be a symplectic vector bundle and suppose that L, L' are lagrangian subbundles such that L_x is transverse to L'_x for each $x \in B$. Then there exists a compatible complex structure J on E satisfying $JL = L'$.*

Proof. Let J_0 be any compatible complex structure on E . Since L' and $J_0 L$ are both transverse to L , we can find a symplectomorphism $T: L \oplus L' \rightarrow L \oplus J_0 L$ which preserves the subbundle L and maps L' to $J_0 L$. A simple check of the definition then shows that $J = T^{-1} J_0 T$ is a compatible complex structure on E which satisfies $JL = L'$.

□

Example 3.21 If E is a symplectic vector bundle over M , then any pair L, L' of lagrangian subbundles of E define a cohomology class $\mu(L, L') \in H^1(M; \mathbb{Z})$ as follows.

Assuming first that E admits a symplectic trivialization $f: E \rightarrow M \times V$ for some symplectic vector space V , we denote by $f_L, f_{L'}: M \rightarrow \mathcal{L}(V)$ the maps induced by the lagrangian subbundles $f(L), f(L')$ of $M \times V$. Then

$$\mu(L, L') = (f_L^* - f_{L'}^*) \mu_V,$$

where $\mu_V \in H^1(\mathcal{L}(V); \mathbb{Z})$ is the universal Maslov class. From Example 3.11 it follows that this class is independent of the choice of trivialization f .

For nontrivial E , we note that since the symplectic group $Sp(V)$ is connected, it follows that for any loop $\gamma: S^1 \rightarrow M$, the pull-back bundle $\gamma^* E$ is trivial. Thus $\mu(L, L')$ is well-defined by the requirement that for every smooth loop γ in M ,

$$\gamma^* \mu(L, L') = \mu(\gamma^* L, \gamma^* L').$$

△

Example 3.22 As a particular case of Example 3.21, we consider the symplectic *manifold* $\mathbb{R}^{2n} \simeq T^*\mathbb{R}^n$ with its standard symplectic structure. Then the tangent bundle $T(\mathbb{R}^{2n})$ is a symplectic vector bundle over \mathbb{R}^{2n} with a natural “vertical” lagrangian subbundle $V\mathbb{R}^n$ defined as the kernel of π_* , where $\pi: T^*\mathbb{R}^n \rightarrow \mathbb{R}^n$ is the natural projection.

If $\iota: L \rightarrow \mathbb{R}^{2n}$ is a lagrangian immersion, then the symplectic vector bundle $\iota^*T(T^*(\mathbb{R}^n))$ has two lagrangian subbundles, the image L_1 of $\iota_*: TL \rightarrow \iota^*T(T^*(\mathbb{R}^n))$ and $L_2 = \iota^*V\mathbb{R}^n$. The class $\mu_{L,\iota} = \mu(L_1, L_2) \in H^1(L; \mathbb{Z})$ is called the **Maslov class** of (L, ι) .

A check of these definitions shows that the Maslov class of (L, ι) equals the pull-back of the universal Maslov class $\mu_n \in H^1(\mathcal{L}(\mathbb{R}^{2n}); \mathbb{Z})$ by the Gauss map $G: L \rightarrow \mathcal{L}(\mathbb{R}^{2n})$ defined by $G(p) = \iota_*T_pL \subset \mathbb{R}^{2n}$. (See [3] for an interpretation of the Maslov class of a loop γ in L as an intersection index of the loop $G \circ \gamma$ with a singular subvariety in the lagrangian grassmannian).

△

3.2 Cotangent bundles

The cotangent bundle T^*M of any smooth manifold M is equipped with a natural 1-form, known as the **Liouville form**, defined by the formula

$$\alpha_M((x, b))(v) = b(\pi_*v),$$

where $\pi: T^*M \rightarrow M$ is the canonical projection. In local coordinates (x_1, \dots, x_n) on M and corresponding coordinates $(q_1, \dots, q_n, p_1, \dots, p_n)$ on T^*M , the equations

$$q_j(x, b) = x_j(x) \quad p_j(x, b) = b\left(\frac{\partial}{\partial x_j}\right)$$

imply that

$$\alpha_M = \sum_{j=1}^n p_j dq_j.$$

Thus, $-d\alpha_M = \sum_{j=1}^n dq_j \wedge dp_j$ in these coordinates, from which it follows that the form $\omega_M = -d\alpha_M$ is a symplectic structure on T^*M . Note that if $M = \mathbb{R}^n$, then ω_M is just the symplectic structure ω_n on $T^*\mathbb{R}^n \simeq \mathbb{R}^{2n}$ discussed previously, and $\alpha_M = \alpha_n$.

Lagrangian immersions and the Liouville class

Given a lagrangian immersion $\iota: L \rightarrow T^*M$, we set $\pi_L = \pi \circ \iota$, where $\pi: T^*M \rightarrow M$ is the natural projection. Critical points and critical values of π_L are called respectively **singular points** and **caustic points** of L . Finally, (L, ι) is said to be **projectable** if π_L is a diffeomorphism. A nice property of the Liouville 1-form is that it can be used to parametrize the set of projectable lagrangian submanifolds. To do this, we use the notation ι_φ to denote a 1-form φ on M when we want to think of it as a map from M to T^*M .

Lemma 3.23 *Let $\varphi \in \Omega^1(M)$. Then*

$$\iota_\varphi^* \alpha_M = \varphi.$$

Proof. Because ι_φ is a section of T^*M , it satisfies $\pi \circ \iota_\varphi = \text{id}_M$. By the definition of α_M , it follows that for each $v \in T_p M$,

$$\iota_\varphi^* \alpha_M(p)(v) = \alpha_M(\iota_\varphi(p))(\iota_{\varphi*} v) = \langle \iota_\varphi(p), \pi_*(\iota_{\varphi*} v) \rangle = \langle \iota_\varphi(p), v \rangle.$$

□

For this reason, α_M is often described as the “tautological” 1-form on T^*M . Taking exterior derivatives on both sides of the equation in Lemma 3.23, we get

$$d\varphi = d\iota_\varphi^* \alpha_M = \iota_\varphi^* d\alpha_M = -\iota_\varphi^* \omega_M.$$

From this equation we see that the image of φ is a lagrangian submanifold of T^*M precisely when the form φ is closed. This proves

Proposition 3.24 *The relation $\varphi \leftrightarrow (M, \iota_\varphi)$ defines a natural bijective correspondence between the the vector space of closed 1-forms on M and the set of projectable lagrangian submanifolds of T^*M .*

Generalizing our WKB terminology, we will call $S : M \rightarrow \mathbb{R}$ a **phase function** for a projectable lagrangian embedding $(L, \iota) \subset T^*M$ provided that $\iota(L) = dS(M)$. The preceding remarks imply a simple link between phase functions and the Liouville form:

Lemma 3.25 *If $(L, \iota) \subset T^*M$ is a projectable lagrangian embedding, then $S : M \rightarrow \mathbb{R}$ is a phase function for L if and only if $d(S \circ \pi_L \circ \iota) = \iota^* \alpha_M$.*

Thus, L is the image of an exact 1-form on M if and only if the restriction of the Liouville form to L is itself exact. This motivates the following definition.

Definition 3.26 *If L, M are n -manifolds and $\iota : L \rightarrow T^*M$ is an immersion such that $\iota^* \alpha_M$ is exact, then ι is called an **exact lagrangian immersion**.*

If $\iota : L \rightarrow T^*M$ is an exact lagrangian immersion, then Lemma 3.25 suggests that the primitive of $\iota^* \alpha_M$ is a sort of generalized phase function for (L, ι) which lives on the manifold L itself. We will return to this important viewpoint in the next chapter.

Example 3.27 A simple application of Stokes’ theorem shows that an embedded circle in the phase plane cannot be exact, although it is the image of an exact lagrangian *immersion* of \mathbb{R} .

A general class of exact lagrangian submanifolds can be identified as follows. Associated to a smooth submanifold $N \subset M$ is the submanifold

$$N^\perp = \{(x, p) \in T^*M : x \in N, T_x N \subset \ker(p)\},$$

known as the **conormal bundle** to N . From this definition it follows easily that $\dim T^*M = 2 \dim N^\perp$, while the Liouville form of T^*M vanishes on N^\perp for any N .

If \mathcal{F} is a smooth foliation of M , then the union of the conormal bundles to the leaves of \mathcal{F} is a smooth submanifold of T^*M foliated by lagrangian submanifolds and is thus coisotropic (see Example 3.19).

△

Although many lagrangian immersions $\iota : L \rightarrow T^*M$ are not exact, the form $\iota^*\alpha_M$ is always closed, since $d\iota^*\alpha_M = \iota^*\omega_M = 0$. The deRham cohomology class $\lambda_{L,\iota} \in H^1(L; \mathbb{R})$ induced by this form will play an important role in the quantization procedures of the next chapter and is known as the **Liouville class** of (L, ι) .

Example 3.28 To generalize the picture described in Example 3.27, we consider a smooth manifold M , together with a submanifold $N \subset M$ and a closed 1-form β on N . Then

$$N_\beta^\perp = \{(x, p) \in T^*M : x \in N \text{ } p|_{T_x N} = \beta(x)\}$$

is a lagrangian submanifold of T^*M whose Liouville class equals $[\pi_N^*\beta] \in H^1(N_\beta^\perp; \mathbb{R})$, where $\pi_N : N_\beta^\perp \rightarrow N$ is here the restriction of the natural projection $\pi : T^*M \rightarrow M$.

△

Fiber-preserving symplectomorphisms

On each fiber of the projection $\pi : T^*M \rightarrow M$, the pull-back of α_M vanishes, so the fibers are lagrangian submanifolds. Thus, the vertical bundle $VM = \ker \pi_*$ is a lagrangian subbundle of $T(T^*M)$. Since α_M vanishes on the zero section $Z_M \subset T^*M$, it follows that Z_M is lagrangian as well, and the subbundles TZ_M and VM define a canonical lagrangian splitting of $T(T^*M)$ over Z_M .

A 1-form β on M defines a diffeomorphism f_β of T^*M by fiber-wise affine translation

$$f_\beta(x, p) = (x, p + \beta(x)).$$

It is easy to see that this map satisfies

$$f_\beta^*\alpha_M = \alpha_M + \pi^*\beta,$$

so f_β is a symplectomorphism of T^*M if and only if β is closed.

Theorem 3.29 *If a symplectomorphism $f : T^*M \rightarrow T^*M$ preserves each fiber of the projection $\pi : T^*M \rightarrow M$, then $f = f_\beta$ for a closed 1-form β on M .*

Proof. Fix a point $(x_0, p_0) \in T^*M$ and let ψ be a closed 1-form on M such that $\psi(x_0) = (x_0, p_0)$. Since f is symplectic, the form $\mu = f \circ \psi$ is also closed, and thus the map

$$h = f_\mu^{-1} \circ f \circ f_\psi$$

is a symplectomorphism of T^*M which preserves fibers and fixes the zero section $Z_M \subset T^*M$. Moreover, since the derivative Dh preserves the lagrangian splitting of $T(T^*M)$ along Z_M and equals the identity on TZ_M , we can conclude from Example 3.1 that Dh is the identity at all points of Z_M . Consequently, the fiber-derivative of f at the arbitrary point (x_0, p_0) equals the identity, so f is a translation on each fiber. Defining $\beta(x) = f(x, 0)$, we have $f = f_\beta$.

□

If β is a closed 1-form on M , then the flow f_t of the vector field $X_\beta = -\tilde{\omega}_M^{-1}(\pi^*\beta)$ is symplectic; since $VM \subset \ker \pi^*\beta$, the Hamilton-Jacobi theorem implies furthermore that the flow f_t satisfies the hypotheses of Theorem 3.29.

Corollary 3.30 *For any closed 1-form β on M , the time-1 map $f = f_1$ of the flow of X_β equals f_β .*

Proof. By Theorem 3.29, the assertion will follow provided that we can show that $f^*\alpha_M = \alpha_M + \pi^*\beta$. To this end, note that the definition of the Lie derivative shows that f satisfies

$$f^*\alpha_M = \alpha_M + \int_0^1 \frac{d}{dt}(f_t^*\alpha_M) dt = \alpha_M + \int_0^1 f_t^*(\mathcal{L}_{X_\beta}\alpha_M) dt.$$

By Cartan's formula for the Lie derivative, we have

$$\mathcal{L}_{X_\beta}\alpha_M = d(X_\beta \lrcorner \alpha_M) - X_\beta \lrcorner \omega_M = \pi^*\beta,$$

the latter equality following from the fact that $X_\beta \subset VM \subset \ker \alpha_M$ and $d\alpha_M = -\omega_M$. Another application of Cartan's formula, combined with the assumption that β is closed shows that

$$\mathcal{L}_{X_\beta}\pi^*\beta = 0,$$

and so $f_t^*\pi^*\beta = \pi^*\beta$ for all t . Inserting these computations into the expression for $f^*\alpha_M$ above, we obtain

$$f^*\alpha_M = \alpha_M + \pi^*\beta.$$

□

Using Theorem 3.29, we can furthermore classify all fiber-preserving symplectomorphisms from T^*M to T^*N .

Corollary 3.31 *Any fiber-preserving symplectomorphism $F: T^*M \rightarrow T^*N$ can be realized as the composition of a fiber-translation in T^*M with the cotangent lift of a diffeomorphism $N \rightarrow M$.*

Proof. By composing F with a fiber-translation in T^*M we may assume that F maps the zero section of T^*M to that of T^*N . The restriction of F^{-1} to the zero sections then induces a diffeomorphism $f: N \rightarrow M$ such that the composition $F \circ (f^{-1})^*$ is a fiber-preserving symplectomorphism of T^*N which fixes the zero section. From the preceding theorem, we conclude that $F = f^*$.

□

The Schwartz transform

If M, N are smooth manifolds, then the map $S_{M,N} : \overline{T^*M} \times T^*N \rightarrow T^*(M \times N)$ defined in local coordinates by

$$((x, \xi), (y, \eta)) \mapsto (x, y, -\xi, \eta)$$

is a symplectomorphism which we will call the **Schwartz transform**.⁵ An elementary, but fundamental property of this mapping can be described as follows.

Proposition 3.32 *If M, N are smooth manifolds, then the Schwartz transform $S_{M,N}$ satisfies*

$$(S_{M,N})^* \alpha_{M \times N} = \alpha_M \oplus -\alpha_N.$$

In particular, $S_{M,N}$ induces a diffeomorphism of zero sections

$$Z_M \times Z_N \simeq Z_{M \times N}$$

and an isomorphism of vertical bundles

$$VM \oplus VN \simeq V(M \times N).$$

Using the Schwartz transform, we associate to any symplectomorphism $F : T^*M \rightarrow T^*N$ the lagrangian embedding $\iota_F : T^*M \rightarrow T^*(M \times N)$ defined as the composition of $S_{M,N}$ with the graph $\Gamma_F : T^*M \rightarrow \overline{T^*M} \times T^*N$.

Example 3.33 By Corollary 3.31, a fiber-preserving symplectomorphism $F : T^*M \rightarrow T^*N$ equals the composition of fiber-wise translation by a closed 1-form β on M with the cotangent lift of a diffeomorphism $g : N \rightarrow M$. A computation shows that if $\Gamma \subset M \times N$ is the graph of g and $p : \Gamma \rightarrow M$ is the natural projection, then the image of the composition of the lagrangian embedding (T^*M, ι_F) with the Schwartz transform $S_{M,N}$ equals the submanifold $\Gamma_{p^*\beta}^\perp \subset T^*(M \times N)$ defined in Example 3.28. In particular, if F is the cotangent lift of g , then the image of (T^*M, ι_F) equals the conormal bundle of Γ .

△

Finally, we note that multiplying the cotangent vectors in T^*M by -1 defines a symplectomorphism $T^*M \rightarrow \overline{T^*M}$ which can be combined with the Schwartz transform $S_{M,N}$ to arrive at the usual symplectomorphism $T^*M \times T^*N \simeq T^*(M \times N)$. Thus in the special case of cotangent bundles, dualizing and taking products leads to nothing new.

3.3 Mechanics on manifolds

With the techniques of symplectic geometry at our disposal, we are ready to extend our discussion of mechanics to more general configuration spaces. Our description begins with a comparison of the classical and quantum viewpoints, largely paralleling the earlier material on the 1-dimensional harmonic oscillator given in the introduction. We then turn to the semi-classical approximation and its geometric counterpart in this new context, setting the stage for the quantization problem in the next chapter.

⁵The name comes from the relation of this construction to the Schwartz kernels of operators (see Section 6.2).

The classical picture

The hamiltonian description of classical motions in a configuration space M begins with the classical phase space T^*M . A riemannian metric $g = (g_{ij})$ on M induces an inner product on the fibers of the cotangent bundle T^*M , and a “kinetic energy” function which in local coordinates (q, p) is given by

$$k_M(q, p) = \frac{1}{2} \sum_{i,j} g^{ij}(q) p_i p_j,$$

where g^{ij} is the inverse matrix to g_{ij} . Regular level sets of k_M are sphere bundles over M . The hamiltonian flow associated to k_M is called the **co-geodesic flow** due to its relation with the riemannian structure of M described in the following theorem.

Theorem 3.34 *If M is a riemannian manifold, then integral curves of the co-geodesic flow project via π to geodesics on M .*

A proof of this theorem can be given in local coordinates by using Hamilton’s equations

$$\dot{q}_i = \frac{\partial H}{\partial p_i} = \sum_j g^{ij} p_j \quad \dot{p}_i = -\frac{\partial H}{\partial q_i} = -\frac{1}{2} \sum_{u,v} \frac{\partial g^{uv}}{\partial q_i} p_u p_v$$

to derive the geodesic equation

$$\ddot{q}_k + \sum_{i,j} \Gamma_{ij}^k \dot{q}_i \dot{q}_j = 0.$$

For details, see [36].

In physical terms, Theorem 3.34 states that a free particle on a manifold must move along a geodesic. A smooth, real-valued potential $V : M \rightarrow \mathbb{R}$ induces the hamiltonian function

$$H(q, p) = k_M(q, p) + V(q)$$

on T^*M . Integral curves of the hamiltonian flow of H then project to classical trajectories of a particle on M subject to the potential V .

The quantum mechanical picture

For the time being, we will assume that the Schrödinger operator on a riemannian manifold M with potential function V is defined in analogy with the flat case of \mathbb{R}^n with its standard metric. That is, we first define the operator on the function space $C^\infty(M)$ by

$$\hat{H} = -\frac{\hbar^2}{2m} \Delta + m_V,$$

where Δ denotes the Laplace-Beltrami operator. As before, \hat{H} induces a (densely defined) operator \hat{H} on the intrinsic Hilbert space \mathfrak{H}_M of M by the equation

$$\hat{H}(a|dx|^{1/2}) = (\hat{H}a)|dx|^{1/2},$$

where $|dx|$ is the natural density associated to the metric on M , and the time-independent Schrödinger equation on M assumes the familiar form

$$(\hat{H} - E)\varphi = 0.$$

The advantage of this viewpoint is that both the classical state space T^*M and the quantum state space \mathfrak{H}_M are objects intrinsically associated to the underlying differential manifold M . The dynamics on both objects are determined by the choice of metric on M .

The semi-classical approximation

The basic WKB technique for constructing semi-classical solutions to the Schrödinger equation on M proceeds as in Section 2.2. Specifically, a half-density of the form $e^{iS/\hbar}a$ is a second-order approximate solution of the eigenvector problem $(\hat{H} - E)\varphi = 0$ provided that the phase function $S: M \rightarrow \mathbb{R}$ satisfies the Hamilton-Jacobi equation

$$H \circ \iota_{dS} = E,$$

and the half-density a satisfies the homogeneous transport equation, which assumes the coordinate-free form

$$a\Delta S + 2\mathcal{L}_{\nabla S}a = 0.$$

We can formulate this construction abstractly by considering first a projectable, exact lagrangian embedding $\iota: L \rightarrow T^*M$. By definition, this means that $\pi_L = \pi \circ \iota$ is a diffeomorphism, where $\pi: T^*M \rightarrow M$ is the natural projection, and Lemma 3.25 implies that for any primitive $\phi: L \rightarrow \mathbb{R}$ of $\iota^*\alpha_M$, the composition $S = \phi \circ \pi_L^{-1}$ is a phase function for (L, ι) .

Now if $H: T^*M \rightarrow \mathbb{R}$ is any smooth function, then (L, ι) satisfies the Hamilton-Jacobi equation provided that E is a regular value of H and

$$H \circ \iota = E.$$

In this case, the embedding ι and hamiltonian vector field X_H of H induce a nonsingular vector field $X_{H,\iota}$ on L (see Example 3.13). If a is a half-density on L , then the requirement that $(\pi_L^{-1})^*a$ satisfy the homogeneous transport equation on M becomes

$$\mathcal{L}_{X_{H,\iota}}a = 0.$$

If these conditions are satisfied, then the half-density $e^{i\phi/\hbar}a$ on L can be quantized (i.e. pulled-back) to yield a second-order approximate solution $(\pi_L^{-1})^*e^{i\phi/\hbar}a$ to the Schrödinger equation on M , as above.

This interpretation of the WKB approximation leads us to consider a semi-classical state as a quadruple (L, ι, ϕ, a) comprised of a projectable, exact lagrangian embedding $\iota: L \rightarrow T^*M$, a “generalized” phase function $\phi: L \rightarrow \mathbb{R}$ satisfying $d\phi = \iota^*\alpha_M$, and a half-density a on L . Our correspondence table now assumes the form

Object	Classical version	Quantum version
basic space	T^*M	\mathfrak{H}_M
state	(L, ι, ϕ, a) as above	on M
time-evolution	Hamilton's equations	Schrödinger equation
generator of evolution	function H on T^*M	operator \hat{H} on \mathfrak{H}_M
stationary state	state (L, ι, ϕ, a) such that $H \circ \iota = E$ and $\mathcal{L}_{X_H, \iota} a = 0$	eigenvector of \hat{H}

Since the Hamilton-Jacobi and transport equations above make sense for any triple (L, ι, a) consisting of an arbitrary lagrangian immersion $\iota: L \rightarrow T^*M$ and a half-density a on L , it is tempting to regard (L, ι, a) as a further generalization of the concept of semi-classical state in T^*M , in which we drop the conditions of projectability and exactness. Our goal in the next chapter will be to determine when and how such “geometric” semi-classical states can be used to construct “analytical” semi-classical approximate solutions to the time-independent Schrödinger equation.

4 Quantization in Cotangent Bundles

This chapter deals with the problem of constructing semi-classical approximate solutions to the time-independent Schrödinger equation from the data contained in a “geometric” semi-classical state. The starting point will be a lagrangian immersion $\iota: L \rightarrow T^*M$ whose image is contained in a regular level set of the classical hamiltonian H associated to some metric and potential on the configuration space M . Given a half-density a on L which is invariant under the flow induced by H , our goal is to use this data in order to construct, in a more or less systematic way, an approximate solution to Schrödinger’s equation on M .

This process will be referred to as **quantization**. (It is, of course, only one of many operations which go by this name.) As we shall see, the question of whether or not a given semi-classical state (L, ι, a) can be quantized is answered largely in terms of the geometry of the lagrangian immersion (L, ι) . For this reason, we will often speak loosely of “quantizing lagrangian submanifolds” or “quantizable lagrangian submanifolds.”

4.1 Prequantization

The simplest quantization procedure consists of pulling-back half-densities from projectable lagrangian embeddings in T^*M . As seen in Chapter 3, a triple (L, ι, a) for which $\iota: L \rightarrow T^*M$ is a projectable *exact* lagrangian embedding is quantized by choosing a primitive ϕ of $\iota^*\alpha_M$ and forming the half-density

$$I_{\hbar}(L, \iota, a) \stackrel{def}{=} (\pi_L^{-1})^* e^{i\phi/\hbar} a$$

on M . The choice of ϕ is unique only up to an additive constant, which leads to an ambiguity in the overall phase of $I_{\hbar}(L, \iota, a)$. (This ambiguity was overcome in Chapter 3 by including the choice of ϕ in the definition of a semi-classical state).

To generalize this procedure, suppose now that $\iota: L \rightarrow T^*M$ is a projectable, but not necessarily exact lagrangian embedding. Since the 1-form $\iota^*\alpha_M$ on L is closed, it is locally exact by the Poincaré lemma, and so we can choose a good cover $\{L_j\}$ of L (see Appendix C) and functions $\phi_j: L_j \rightarrow \mathbb{R}$ such that $d\phi_j = \iota^*\alpha_M|_{L_j}$. Given a half-density a on L , we set $a_j = a|_{L_j}$ and define a half-density on $\pi_L(L_j)$ by quantizing $(L_j, \iota|_{L_j}, \phi_j, a_j)$ in the sense above:

$$I_j = (\pi_{L_j}^{-1})^* e^{i\phi_j/\hbar} a_j.$$

To quantize (L, ι, a) , we must piece together the I_j to form a well-defined global half-density $I_{\hbar}(L, \iota, a)$ on M . This is possible for arbitrary a provided that the functions ϕ_j can be chosen so that the oscillatory coefficients $e^{i\phi_j/\hbar}$ agree where their domains overlap; that is, we must have

$$\phi_j - \phi_k \in \mathbb{Z}_{\hbar} \stackrel{def}{=} 2\pi\hbar \cdot \mathbb{Z}$$

on each $L_j \cap L_k$. According to the discussion in Appendix C, this is precisely the condition that the Liouville class $\lambda_{L, \iota}$ be \hbar -integral. At this point, however, we are forced to confront once more the nature of \hbar . If it is a formal variable, the notion of \hbar -integrality is meaningless. If \hbar denotes a number ranging over an interval of the real numbers, then the Liouville class will be $2\pi\hbar$ times an integral class for all \hbar only if it is zero, i.e. only for an exact lagrangian

submanifold. Since we are specifically trying to go beyond the exact case, this interpretation is not acceptable either. Instead, we will make the following compromise.

Definition 4.1 *A projectable lagrangian submanifold $(L, \iota) \subset T^*M$ is **quantizable** if its Liouville class $\lambda_{L, \iota}$ is \hbar -integral for some $\hbar \in \mathbb{R}_+$. The values of \hbar for which this condition holds will be called **admissible** for (L, ι) .*

If L is quantizable but not exact, the set of all admissible \hbar forms a sequence converging to zero, consisting of the numbers \hbar_0/k , where \hbar_0 is the largest such number, and k runs over the positive integers. If L is exact, all $\hbar > 0$ are admissible.

Example 4.2 Let $M = S^1$ and consider the closed 1-form $\beta = p d\theta$ on S^1 , for $p \in \mathbb{R}$. The cohomology class represented by this form is \hbar -integral provided that

$$p \int_{S^1} d\theta \in \mathbb{Z}_{\hbar},$$

and so the admissible values of \hbar are the numbers $\{p/k : k \in \mathbb{Z}_+\}$. Geometrically this means that all horizontal circles in the cylinder T^*S^1 are quantizable in the sense defined above, but with differing sets of admissible \hbar .

The situation changes as soon as we consider the closed 1-form $\tau = a d\theta_1 + b d\theta_2$ on the torus $S^1 \times S^1$. If τ is \hbar -integral for some $\hbar \in \mathbb{R}_+$, then a/\hbar and b/\hbar are both integers, meaning that a/b is a rational number if $b \neq 0$. In general, the condition that closed 1-form β on a manifold M be quantizable is equivalent to the requirement that the ratio of any two nonzero periods of β be rational. (A *period* of a closed 1-form β on M is any number obtained by integrating β around some closed loop in M). Thus, the class of projectable quantizable lagrangian submanifolds of T^*M is rather limited whenever $\dim(H^1(M; \mathbb{R})) > 1$.

△

The simple quantization technique described above does not generalize immediately to non-projectable immersed lagrangian submanifolds $(L, \iota) \subset T^*M$, since π_L cannot be used to push-forward half-densities from L to M . For the time being, however, we will focus on the set of regular points of π_L in order to pass from half-densities on L to half-densities defined near non-caustic points of L . Regardless of how this is to be carried out, it is desirable that quantization be linear with respect to half-densities:

$$I_{\hbar}(L, \iota, sa_1 + a_2) = s \cdot I_{\hbar}(L, \iota, a_1) + I_{\hbar}(L, \iota, a_2).$$

Certainly this condition holds for the procedure we have been using in the projectable case, and it is reasonable to adopt as a general rule. One consequence is

$$I_{\hbar}(L, \iota, 0) = 0.$$

If a semi-classical state is represented by the union of a disjoint pair $(L_1, \iota_1, a_1), (L_2, \iota_2, a_2)$ of lagrangian submanifolds carrying half-densities, then by linearity we should have

$$\begin{aligned} I_{\hbar}(L_1 \cup L_2, \iota_1 \cup \iota_2, (a_1, a_2)) &= I_{\hbar}(L_1 \cup L_2, \iota_1 \cup \iota_2, (a_1, 0)) + I_{\hbar}(L_1 \cup L_2, \iota_1 \cup \iota_2, (0, a_2)) \\ &= I_{\hbar}(L_1, \iota_1, a_1) + I_{\hbar}(L_2, \iota_2, a_2). \end{aligned}$$

Now consider an arbitrary immersed lagrangian submanifold $(L, \iota) \subset T^*M$ and half-density a on L . If $p \in \pi_L(L)$ is non-caustic and π_L is proper, then there is a contractible neighborhood $U \subset M$ of p for which $\pi_L^{-1}(U)$ consists of finitely many disjoint open subsets $L_j \subset L$ such that each $(L_j, \iota|_{L_j})$ is a projectable lagrangian submanifold of T^*U . Choosing a generalized phase function $\phi_j: L_j \rightarrow \mathbb{R}$ for each L_j , we note that by the preceding remarks, the quantization of (L, ι, a) should look something like

$$\sum_j (\pi_{L_j}^{-1})^* e^{i\phi_j/\hbar} a$$

on U . As before, the requirement that ϕ_j be a generalized phase function for L_j determines each ϕ_j only up to an additive constant, and so the meaning of the preceding sum is ambiguous. To quantize half-densities on L consistently, we must therefore decide how to specify the relative phases of the oscillatory coefficients $e^{i\phi_j/\hbar}$.

If L is exact, then we may use any function $\phi: L \rightarrow \mathbb{R}$ satisfying $d\phi = \iota^* \alpha_M$ to fix phases, as in the following example.

Example 4.3 For $L = \mathbb{R}$, the lagrangian embedding $\iota: L \rightarrow T^*\mathbb{R}$ given by $\iota(x) = (x^2, x)$ has a singular point at $x = 0$, and we denote by L_+, L_- the right and left projectable components of L , respectively (these correspond to the upper and lower components of the parabola $\iota(L)$). A phase function $\phi: L \rightarrow \mathbb{R}$ for (L, ι) is given by

$$\phi(x) = 2x^3/3.$$

If $a = B(x) |dx|^{1/2}$ is any half-density on L , then since

$$(\pi_L^{-1})^* \frac{\partial}{\partial q} = \pm 2^{-1} q^{-1/2} \frac{\partial}{\partial x},$$

the transformation rule for half-densities implies

$$(\pi_{L_+}^{-1})^* a = 2^{-1/2} q^{-1/4} B(q^{1/2}) |dq|^{1/2} \quad (\pi_{L_-}^{-1})^* a = 2^{-1/2} q^{-1/4} B(-q^{1/2}) |dq|^{1/2}.$$

Thus, the prequantization of (L, ι, a) is given for $q > 0$ by

$$I_{\hbar}(L, \iota, a)(q) = \left(e^{2iq^3/2/3\hbar} B(q^{1/2}) + e^{-2iq^3/2/3\hbar} B(-q^{1/2}) \right) 2^{-1/2} q^{-1/4} |dq|^{1/2}.$$

The parabola $\iota(L)$ lies in the regular level set $H^{-1}(0)$ of the hamiltonian for a constant force field

$$H(q, p) = \frac{1}{2}(p^2 - q),$$

and it is easy to check that the induced vector field $X_{H, \iota}$ on L equals $X_{H, \iota} = (1/2) \partial/\partial x$. Thus, a half-density a on L is invariant under the flow of $X_{H, \iota}$ if and only if $a = B |dx|^{1/2}$ for some $B \in \mathbb{R}$. From the expression above, we obtain

$$I_{\hbar}(L, \iota, a) = \left(e^{2iq^3/2/3\hbar} + e^{-2iq^3/2/3\hbar} \right) 2^{-1/2} q^{1/4} B |dq|^{1/2}$$

as a semi-classical approximate solution to the Schrödinger equation

$$-\frac{\hbar^2}{2} \frac{\partial^2 \psi}{\partial x^2} - \frac{x}{2} \psi = E \psi.$$

Unfortunately, this solution blows up at $q = 0$ (and is not defined for $q < 0$), so we have more work ahead of us. In particular, we have no way as yet to check that letting ϕ be continuous at 0 as a function on L is the right way to assure that we have a good approximation in the immediate vicinity of $q = 0$. In fact, we will see later that this is the wrong choice!

△

Since exactness is only used to insure that the function $e^{i\phi/\hbar}$ is well-defined on L , we can treat certain non-exact cases in a similar way. For this purpose, we make the following provisional definition, generalizing Definition 4.1.

Definition 4.4 *An immersed lagrangian submanifold $(L, \iota) \subset T^*M$ is said to be **prequantizable** if its Liouville class $\lambda_{L, \iota}$ is \hbar -integral for some $\hbar \in \mathbb{R}_+$. The values of \hbar for which this condition holds will again be called **admissible** for (L, ι) .*

If \hbar is admissible for some prequantizable lagrangian immersion (L, ι) , then there exists a good cover $\{V_j\}$ of the manifold L and functions $\phi_j : V_j \rightarrow \mathbb{R}$ such that $d\phi_j = \iota^* \alpha_M|_{V_j}$ and $\phi_j - \phi_k \in \mathbb{Z}\hbar$ on each $V_j \cap V_k$. Consequently, the ϕ_j describe a single function $\phi : L \rightarrow \mathbb{T}_\hbar = \mathbb{R}/\mathbb{Z}\hbar$ which satisfies $d\phi = \iota^* \alpha_M$ and which defines a global oscillatory function $e^{i\phi/\hbar}$ on L . If a is a half-density on L , we can now quantize (L, ι, ϕ, a) by summing the pull-backs of $e^{i\phi/\hbar} a$ to M . In the previous notation, the value of $I_\hbar(L, \iota, \phi, a)$ on U is defined as

$$\sum_j (\pi_{V_j}^{-1})^* e^{i\phi/\hbar} a.$$

Example 4.5 Consider the 1-dimensional harmonic oscillator with hamiltonian

$$H(q, p) = \frac{1}{2}(q^2 + p^2).$$

According to Definition 4.4, a number $\hbar \in \mathbb{R}_+$ is admissible for the level set $H^{-1}(E)$ provided that

$$\int_{H^{-1}(E)} \alpha_1 = 2\pi E \in \mathbb{Z}\hbar.$$

Thus, energy levels of the 1-dimensional harmonic oscillator corresponding to level sets for which a particular value of \hbar is admissible are given by $E = n\hbar$. As one can read in any textbook on quantum mechanics, the actual quantum energy levels are $E = (n + 1/2)\hbar$. The additional 1/2 can be explained geometrically in terms of the non-projectability of the classical energy level curves, as we shall soon see.

△

Prequantum bundles and contact manifolds

Prequantizability can be described geometrically in terms of principal \mathbb{T}_{\hbar} bundles with connection over T^*M . It is customary to make the following definition.

Definition 4.6 *For fixed $\hbar \in \mathbb{R}_+$, the **prequantum \mathbb{T}_{\hbar} bundle** associated to a cotangent bundle (T^*M, ω_M) consists of the trivial principal bundle $Q_{M, \hbar} = T^*M \times \mathbb{T}_{\hbar}$ together with the connection 1-form $\varphi = -\pi^*\alpha_M + d\sigma$.*

Here, σ denotes the multiple-valued linear variable in \mathbb{T}_{\hbar} and $\pi: Q_{M, \hbar} \rightarrow T^*M$ is the bundle projection. If $\iota: L \rightarrow T^*M$ is any lagrangian immersion, then the curvature of the induced connection on $\iota^*Q_{M, \hbar}$ coincides with $\iota^*\omega_M$ and therefore vanishes. The holonomy of this connection is represented by the mod- \mathbb{Z}_{\hbar} reduction of the Liouville class $\lambda_{L, \iota}$.

From the prequantum standpoint, the basic geometric object representing a classical state is therefore a quadruple (L, ι, a, ϕ) consisting of a lagrangian immersion $\iota: L \rightarrow T^*M$, a half-density a on L , and a parallel lift ϕ of L to the \mathbb{T}_{\hbar} bundle $\iota^*Q_{M, \hbar}$. If we associate to $Q_{M, \hbar}$ the prequantum line bundle $\mathcal{E}_{M, \hbar}$ by means of the representation $x \mapsto e^{-ix/\hbar}$ of \mathbb{T}_{\hbar} in $U(1)$, then ϕ induces a parallel section of $\iota^*\mathcal{E}_{M, \hbar}$ corresponding to the (inverse of the) oscillatory function $e^{i\phi/\hbar}$ on L which appeared in the preceding section. This remark proves the following geometric characterization of prequantizability.

Theorem 4.7 *An immersed lagrangian submanifold $(L, \iota) \subset T^*M$ is prequantizable if and only if there exists a nonzero parallel section over L of the line bundle $\iota^*\mathcal{E}_{M, \hbar}$ for some $\hbar > 0$.*

Prequantum \mathbb{T}_{\hbar} bundles and parallel lifts of lagrangian submanifolds constitute our first examples of the fundamental objects of **contact geometry**, the odd-dimensional counterpart of symplectic geometry. Digressing briefly from quantization, we assemble here a few facts about contact manifolds.

Definition 4.8 *A **contact form** on a $(2n+1)$ -dimensional manifold Q is a 1-form φ such that $\varphi \wedge (d\varphi)^n$ vanishes nowhere on Q . A manifold endowed with a contact form is called a **strict contact manifold**.*

To interpret the condition on φ we note first that the kernel of any nowhere vanishing 1-form φ defines a $2n$ -dimensional distribution in Q . If ξ, η are (local) vector fields lying in this distribution, we have

$$d\varphi(\xi, \eta) = \xi \cdot \varphi(\eta) - \eta \cdot \varphi(\xi) - \varphi([\xi, \eta]) = -\varphi([\xi, \eta]).$$

This says that the distribution is integrable (in the sense of Frobenius) iff $d\varphi$ is zero on $\ker(\varphi)$. The condition $\varphi \wedge (d\varphi)^n \neq 0$, on the other hand, means that the kernel of $d\varphi$ is 1-dimensional and everywhere transverse to $\ker(\varphi)$. Consequently, $d\varphi$ is a linear symplectic form on $\ker(\varphi)$, and the “largest” integral submanifolds of $\ker(\varphi)$ are n -dimensional (i.e. $\ker(\varphi)$ is “maximally non-integrable”).

Definition 4.9 *A **legendrian submanifold** of a $2n+1$ -dimensional strict contact manifold (Q, φ) is an n -dimensional integral submanifold for φ .*

If φ is a contact form and f a nowhere-vanishing function on Q , then $f \cdot \varphi$ is again a contact form which is said to be equivalent to φ , since they have the same legendrian submanifolds. This leads to the following definition.

Definition 4.10 *A contact structure on a manifold M is a codimension one subbundle $E \subset TM$ which is locally defined by contact forms, and a manifold endowed with such a structure is called simply a **contact manifold**.*

When the quotient TM/E is trivial, E is the kernel of a globally defined contact form. Such a contact structure is called **coorientable**.

Our basic example of a strict contact manifold is furnished by a prequantum \mathbb{T}_\hbar bundle $Q_{M,\hbar}$ over a cotangent bundle T^*M ; the image of a parallel lift of a lagrangian immersion $\iota: L \rightarrow T^*M$ to $Q_{M,\hbar}$ is an immersed legendrian submanifold of $Q_{M,\hbar}$. Although we will not pursue the idea in these notes, a possible generalization of the quantization procedure described above might therefore begin by re-interpreting geometric semi-classical states as triples (R, j, a) consisting of a legendrian immersion (R, j) in a contact manifold together with a half-density a on R .

4.2 The Maslov correction

It turns out that the naive quantization procedure of the preceding section is incorrect, since it ignores a certain structure which arises from the relation of L to the fibers of the projection $T^*M \xrightarrow{\pi} M$. This factor will be incorporated into our quantization procedure using a procedure due to Maslov. To begin, we will illustrate this idea in the case of lagrangian submanifolds of the phase plane.

Given a lagrangian immersion $\iota: L \rightarrow T^*\mathbb{R} \simeq \mathbb{R}^2$, we denote by π_p the composition of ι with the projection of \mathbb{R}^2 onto the p -axis. If π_p is a diffeomorphism, then (L, ι) is said to be p -projectable, in which case there exists an “alternate” generalized phase function $\tau: L \rightarrow \mathbb{R}$ satisfying $d\tau = \iota^*(-q dp)$, obtained by thinking of \mathbb{R}^2 as the cotangent bundle of p -space.

A simple example of an embedded lagrangian submanifold of $T^*\mathbb{R}$ which does not project diffeomorphically onto the q -axis is a vertical line, or fiber, of the form $\iota(x) = (q_0, x)$ for $x \in \mathbb{R}$. Since the wave function corresponding to a constant half-density a on L should correspond to a probability distribution describing the position of a particle at q_0 with completely indeterminate momentum, it should be a delta function supported at q_0 . Following an idea of Maslov, we analyze this situation by pretending that p is position and q momentum, and then quantizing to obtain a function on p -space. Using the phase function $\tau(x) = -q_0x$ on L , we obtain

$$(\pi_p^{-1})^* e^{i\tau/\hbar} |dx|^{1/2} = e^{-iq_0p/\hbar} q_0 |dp|^{1/2}.$$

The result is exactly the asymptotic Fourier transform (see Appendix B) of the delta function which we guessed above!

In its simplest form, Maslov’s technique is to suppose that $(L, \iota) \subset T^*\mathbb{R}$ is p -projectable, so that $d\tau = \iota^*(-q dp)$ for some phase function τ on L . If a is a half-density on L , we define a function B on p -space by the equation

$$B |dp|^{1/2} = (\pi_p^{-1})^* e^{i\tau/\hbar} a.$$

The Maslov quantization of (L, ι, τ, a) is then given by the half-density

$$J_{\hbar}(L, \iota, \tau, a) \stackrel{\text{def}}{=} \mathcal{F}_{\hbar}^{-1}(B) |dq|^{1/2}$$

on q -space, where \mathcal{F}_{\hbar} denotes the asymptotic Fourier transform. To relate this procedure to our earlier quantization by pull-back, we must compare the results in the case of lagrangian submanifolds $L \subset \mathbb{R}^2$ which are bi-projectable, i.e. projectable in both the q - and p -directions. For simplicity, we begin with the example of linear lagrangian subspaces.

Example 4.11 For real $k \neq 0$, consider the lagrangian embedding $\iota: \mathbb{R} \rightarrow T^*\mathbb{R}$ given by $\iota(x) = (x, kx)$. Generalized phase functions on (L, ι) for the forms $p dq$ and $-q dp$ are given by $\phi(x) = kx^2/2$ and $\tau(x) = -kx^2/2$, respectively. If a is a constant half-density on L , then the transformation rule for half-densities implies

$$(\pi_L^{-1})^* a = A |dq|^{1/2} \quad (\pi_p^{-1})^* a = |k|^{-1/2} A |dp|^{1/2}$$

for a real constant A determined by a . Quantization by pull-back therefore gives

$$I_{\hbar}(L, \iota, \phi, a) = e^{ikq^2/2\hbar} A |dq|^{1/2}.$$

On the other hand, we have $(\pi_p)^* \tau(p) = -p^2/2k$, and a computation⁶ shows that

$$\mathcal{F}_{\hbar}^{-1}((\pi_p^{-1})^* e^{i\tau/\hbar})(q) = |k|^{1/2} e^{-i\pi \cdot \text{sgn}(k)/4} e^{ikq^2/2\hbar},$$

and so Maslov's technique yields

$$J_{\hbar}(L, \iota, \tau, a) = e^{-i\pi \cdot \text{sgn}(k)/4} I_{\hbar}(L, \iota, \phi, a).$$

Thus, the half-density obtained from (L, ι, a) by Maslov quantization differs from the simple pull-back by a constant phase shift.

△

To establish a similar correspondence in somewhat greater generality, consider an arbitrary bi-projectable lagrangian embedding $(L, \iota) \subset \mathbb{R}^2$ with phase functions ϕ and τ corresponding to $p dq$ and $-q dp$, respectively. For simplicity, we will assume that the additive constants in ϕ and τ are chosen so that

$$\phi = \tau + \iota^*(qp).$$

Next, let $S(q)$ and $T(p)$ be the functions defined on q - and p -space by pull-back:

$$S = \phi \circ \pi_L^{-1} \quad T = \tau \circ \pi_p^{-1},$$

so that S and T satisfy the Legendre transform relation (see [2])

$$S(q) = -p(q) T'(p(q)) + T(p(q)),$$

⁶See, for example, [32, Vol.1, Thm.7.6.1] .

where $p(q) = S'(q)$. From this relation, it follows easily that

$$T''(p(q)) = -(S''(q))^{-1}.$$

A half-density a on L determines functions $A(q)$ and $B(p)$ such that

$$(\pi_L^{-1})^* a = A |dq|^{1/2} \quad (\pi_p^{-1})^* a = B |dp|^{1/2}$$

and

$$A(q) = |S''(q)|^{1/2} B(p(q)).$$

For each q , we must now compare the Maslov half-density

$$J_{\hbar}(L, \iota, \tau, a) = (2\pi\hbar)^{-1/2} \int_{\mathbb{R}} e^{i(pq+T(p))/\hbar} B(p) dp |dq|^{1/2}$$

with that obtained by pull-back:

$$I_{\hbar}(L, \iota, \phi, a) = e^{iS/\hbar} A |dq|^{1/2}.$$

To this end, we set $k(q) = T''(p(q))$ and apply the principle of stationary phase (see Appendix B). The critical point of the exponent $pq + T(p)$ occurs where $q = -T'(p)$, i.e. where $p = S'(q) = p(q)$. Hence,

$$\int_{\mathbb{R}} e^{i(pq+T(p))/\hbar} B(p) dp = (2\pi\hbar)^{1/2} e^{-i\pi \cdot \text{sgn}(k)/4} e^{iS(q)/\hbar} |k(q)|^{-1/2} B(p(q)) + O(\hbar^{3/2}).$$

Thus

$$J_{\hbar}(L, \iota, \tau, a) = e^{-i\pi \cdot \text{sgn}(k)/4} I_{\hbar}(L, \iota, \phi, a) + O(\hbar).$$

For bi-projectable (L, ι) , we therefore conclude that Maslov's technique coincides with quantization by pull-back up to a constant phase factor and terms of order \hbar .

The essential difference between the naive prequantization of Section 4.1 and the Maslov quantization of a p -projectable lagrangian embedding (L, ι) which is not q -projectable lies in the relative phase constants of the summands of $I_{\hbar}(L, \iota, a)$, as illustrated by the following example.

Example 4.12 A phase function associated to $-q dp$ for the lagrangian embedding $\iota(x) = (x^2, x)$ of $L = \mathbb{R}$ into \mathbb{R}^2 is given by $\tau(x) = -x^3/3$. The Maslov quantization of a half-density $a = B(x) |dx|^{1/2}$ on L is thus

$$J_{\hbar}(L, \iota, \tau, a) = \mathcal{F}_{\hbar}^{-1} \left(e^{-ip^3/3\hbar} B(p) \right) |dq|^{1/2},$$

since $(\pi_p^{-1})^* a = B(p) |dp|^{1/2}$. For each $q > 0$, critical points of the function $R(p) = pq - p^3/3$ occur precisely when $q = p^2$, i.e. when $(q, p) \in \iota(L)$, and an application of the principle of stationary phase therefore yields two terms corresponding to the upper and lower halves of L . Specifically, we have

$$J_{\hbar}(L, \iota, \tau, a) = \left(e^{-i\pi/4} e^{2iq^{3/2}/3\hbar} B(-q^{1/2}) + e^{i\pi/4} e^{-2iq^{3/2}/3\hbar} B(q^{1/2}) \right) 2^{-1/2} q^{-1/4} |dq|^{1/2} + O(\hbar).$$

Compare this with the result of Example 4.3. The extra phase factors of $e^{\mp i\pi/4}$ make $J_{\hbar}(L, \iota, \tau, a)$ essentially different from the prequantization of (L, ι, a) . However, while the term of order 0 of $J_{\hbar}(L, \iota, \tau, a)$ is, like $I_{\hbar}(L, \iota, \phi, a)$, singular at the caustic point $q = 0$, the full expression for $J_{\hbar}(L, \iota, \tau, a)$ as an integral is perfectly smooth there, at least if a has compact support. This smoothness at caustics is a clear advantage of Maslov quantization.

△

The relative phase factor in the preceding example can be attributed to the fact that the function $T(p) = -p^3/3$ has an inflection point at $p = 0$. More precisely, since T' is convex, a factor of $e^{i\pi/2}$ arises in passing from the upper to the lower half of the parabola; if T' were concave, the situation would be reversed.

The preceding observation leads us to assign an index to closed, immersed curves in the phase plane. A p -dependent phase function T for $L \subset \mathbb{R}^2$ will have inflection points at precisely those p for which $(T'(p), p)$ is a singular point of L . Moreover, the sign of T'' at nearby points depends only on L and not on the choice of T . With these remarks in mind, suppose that (L, ι) is a closed, immersed curve in \mathbb{R}^2 which is non-degenerate in the sense that if T is a p -dependent phase function for a subset of L , then T' has only non-degenerate critical points. Under this assumption, $\text{sgn}(T'')$ changes by ± 2 in the vicinity of a critical point of T' , and we can assign an index to (L, ι) by summing these changes while traversing L in a prescribed direction. The result is twice an integer known as the **Maslov index** $m_{L, \iota}$ of (L, ι) . (Compare [3]).

Example 4.13 The computation of the Maslov index can be interpreted geometrically if we first observe that the non-degeneracy condition requires L to remain on the same side of a fiber $\pi^{-1}(q)$ near a singular point. Since the index only involves the sign of T'' , it follows from Example 4.12 that the integer 1 should be assigned to a critical point of π_L which is traversed in the $-p$ direction to the right of the fiber. In other words, downward motion to the right of the fiber is positive, while the sign changes if *either* the direction of motion *or* the side of the fiber is reversed, but not if both are.

A circle in the phase plane has a right and a left singular point, both of which are positive according our rule when the circle is traversed counterclockwise. The Maslov index of the circle therefore equals 2. In fact, the same is true of any closed *embedded* curve traversed counterclockwise. On the other hand, a figure-eight has Maslov index zero.

The reader is invited to check that if L is a circle with a fixed orientation $\nu \in H_1(L; \mathbb{Z})$, then $m_{L, \iota}$ equals $\langle \mu_{L, \iota}, \nu \rangle$, where $\mu_{L, \iota}$ is the Maslov class of (L, ι) defined in Example 3.22.

△

Our goal is now to modify the prequantization procedure for arbitrary lagrangian submanifolds of the phase plane by incorporating the Maslov index. The basic idea is as follows. Given an immersed lagrangian submanifold $(L, \iota) \subset \mathbb{R}^2$, we first choose a good cover $\{L_j\}$ of L such that the image of each L_j under ι is either q - or p -projectable and no intersection $L_j \cap L_k$ contains a critical point of π_L . Next, we fix a partition of unity $\{h_j\}$ subordinate

to $\{L_j\}$. To quantize a half-density a on L , we quantize each $(L_j, \iota, a \cdot h_j)$ to obtain a half-density I_j on \mathbb{R} either by pull-back or by Maslov's technique. As before, we would then like to define the quantization of (L, ι, a) as the sum

$$I_{\hbar}(L, \iota, a) = \sum_j I_j.$$

In order to specify the relative phases of the I_j , and to make this definition independent of the choice of cover $\{L_j\}$ and partition of unity, we will require that $I_{\hbar}(L, \iota, a)$ coincide up to order \hbar with the usual quantization by pull-back for any half-density a supported in a projectable subset of L . This condition can be precisely formulated in terms of the Maslov index and Liouville class of (L, ι) as follows.

On an open interval U of non-caustic points, each half-density I_j is the sum of the pull-back of half-densities on each component of $L_j \cap \pi_L^{-1}(U)$. If L_j is quantized by pull-back, this statement is obvious; if Maslov's technique is applied to L_j , it follows from an application of the principle of stationary phase as in Example 4.12. On $L_j \cap \pi_L^{-1}(U)$, these half-densities are of the form

$$\tilde{I}_j = e^{-i\pi s_j/4} e^{i\phi_j/\hbar} a.$$

Here, ϕ is a real-valued function on L_j satisfying $d\phi_j = \iota^* \alpha_1$, while s_j are integers depending only on the component of $L_j \cap \pi_L^{-1}(U)$ in question. (More precisely, s_j is zero if L_j is q -projectable and is quantized by pull-back; otherwise s_j equals $\text{sgn}(T'')$ for a suitable p -dependent phase function). For $I_{\hbar}(L, \iota, a)$ to be well-defined, we must choose the functions ϕ_j so that $\tilde{I}_j = \tilde{I}_k$ on each intersection $L_j \cap L_k$ regardless of the particular half-density a . In other words, we require

$$e^{i(\phi_j - \phi_k)/\hbar} e^{-i\pi(s_j - s_k)/4} = 1$$

at each point of $L_j \cap L_k$. Since $\phi_j - \phi_k$ is constant on $L_j \cap L_k$, we can define a_{jk} as the (constant) value of $(\phi_j - \phi_k) - \pi\hbar(s_j - s_k)/4$ on $L_j \cap L_k$, so that our requirement becomes

$$a_{jk} \in \mathbb{Z}_{\hbar}.$$

Evidently, this condition can be fulfilled on any arc of a curve in the phase plane. If L is a circle, then this condition implies that the sum of any of the a_{jk} lies in \mathbb{Z}_{\hbar} , or, in other words, that the Maslov index $m_{L, \iota}$ of (L, ι) satisfies

$$\frac{\pi\hbar}{2} m_{L, \iota} + \int_L \iota^* \alpha_1 \in \mathbb{Z}_{\hbar}.$$

This is the simplest version of the **Maslov quantization condition**.

Example 4.14 Returning to the harmonic oscillator of Example 4.5, we see that the level set $H^{-1}(E)$ satisfies the Maslov condition provided that for some integer n ,

$$E = (n + 1/2)\hbar.$$

Allowable energy levels in this case therefore correspond to the Bohr-Sommerfeld condition, which actually gives the precise energy levels for the quantum harmonic oscillator.

△

A general quantization scheme

Motivated by the simple results above, our aim in the next sections will be to develop a systematic method for quantizing lagrangian submanifolds of cotangent bundles. The basis of this method will again be Maslov's technique, which relies in this context on the concept of generalized phase functions.

Local parametrizations of an immersed lagrangian submanifold $(L, \iota) \subset T^*M$ defined by such phase functions will first enable us to quantize a given half-density locally on L by means of a slightly more general version of the (inverse) asymptotic Fourier transform. The result will be a collection of half-densities on M . In order for these half-densities to piece together appropriately, it is necessary and sufficient that L satisfy a general version of the Maslov quantization condition, which we formulate in the next section using the Maslov class.

4.3 Phase functions and lagrangian submanifolds

In this section, we generalize the concept of phase functions to non-projectable lagrangian submanifolds of cotangent bundles. Roughly speaking, the idea is the following. As we saw in Chapter 3, a projectable lagrangian submanifold of T^*N can be locally parametrized by the differential of a function f on $U \subset N$, viewed as a mapping $df : U \rightarrow T^*N$. For each $p \in U$, the meaning of df_p as an element of T_p^*N is that for any smooth curve γ in N satisfying $\gamma(0) = p$, we have $\langle df_p, \dot{\gamma}(0) \rangle \stackrel{\text{def}}{=} (f \circ \gamma)'(0)$. To parametrize more general lagrangian submanifolds of T^*N in a similar way, we can begin with a function $\varphi : U \times \mathbb{R}^m \rightarrow \mathbb{R}$ together with a point $\tilde{p} = (p, v) \in U \times \mathbb{R}^m$ and attempt to define $d\varphi_{\tilde{p}} \in T_p^*N$ by $\langle d\varphi_{\tilde{p}}, \dot{\gamma}(0) \rangle = (\varphi \circ \tilde{\gamma})'(0)$ for any lift $\tilde{\gamma}$ of γ to the product $U \times \mathbb{R}^m$ such that $\tilde{\gamma}(0) = \tilde{p}$. In general, this fails, since the value of the directional derivative $(\varphi \circ \tilde{\gamma})'(0)$ depends on the lift $\tilde{\gamma}$. If, however, the fiber-derivative $\partial\varphi/\partial\theta$ vanishes at \tilde{p} , the expression for $d\varphi_{\tilde{p}}$ produces a well-defined element of T_p^*N . As we shall see, the assumption that the map $\partial\varphi/\partial\theta : U \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ is transverse to 0 implies that the *fiber critical set*

$$\Sigma_\varphi = \left\{ \tilde{p} \in U \times \mathbb{R}^m : \frac{\partial\varphi}{\partial\theta} = 0 \right\}$$

is a smooth submanifold of $U \times \mathbb{R}^m$, and the assignment $\tilde{p} \mapsto d\varphi_{\tilde{p}}$ defines a lagrangian immersion of Σ_φ into T^*N . In general, the restriction of the projection $U \times \mathbb{R}^m \rightarrow U$ to Σ_φ is non-injective, and thus the image of Σ_φ is a *non-projectable* lagrangian submanifold. From the point of view of the WKB method, this generalization amounts to replacing Maslov's ansatz

$$(2\pi\hbar)^{-n/2} \int_{\mathbb{R}^n} e^{i\langle(p,q)+T(p)\rangle/\hbar} a(p) dp |dq|^{1/2}$$

for the solution of Schrödinger's equation by the more general form

$$(2\pi\hbar)^{-m/2} \int_{\mathbb{R}^m} e^{i\phi(q,\theta)/\hbar} a(q, \theta) |d\theta| |dq|^{1/2},$$

where θ is an auxiliary variable in \mathbb{R}^m which may have nothing to do with the variable p dual to q . An advantage of this generalization will be to allow a calculus which is more clearly invariant under changes of coordinates, unlike the previous Fourier transform picture, which requires linear structures on p - and q -space.

To begin, we fix some notation and terminology.⁷ Let M, B be smooth manifolds, and let $p_M: B \rightarrow M$ be a smooth submersion. Dualizing the inclusion $E = \ker(p_{M*}) \xrightarrow{\iota} TB$ gives rise to an exact sequence of vector bundles over B

$$0 \leftarrow E^* \xleftarrow{\iota^*} T^*B \leftarrow E^\perp \leftarrow 0$$

where $E^\perp \subset T^*B$ denotes the annihilator of E . The **fiber-derivative** of a function $\phi: B \rightarrow \mathbb{R}$ is the composition $d_\theta\phi = \iota^* \circ d\phi$, and its *fiber critical set* is defined as

$$\Sigma_\phi = (d_\theta\phi)^{-1}Z_{E^*}.$$

(We will denote the zero section of a vector bundle F by Z_F). The function ϕ is said to be **nondegenerate** if its fiber derivative is transverse to Z_{E^*} , in which case Σ_ϕ is a smooth submanifold of B . At points of Σ_ϕ , the section $d_\theta\phi$ has a well-defined intrinsic derivative (see [26]), denoted $\nabla d_\theta\phi$, which induces for nondegenerate ϕ an exact sequence of vector bundles over Σ_ϕ

$$0 \rightarrow T\Sigma_\phi \rightarrow T_{\Sigma_\phi}B \xrightarrow{\nabla d_\theta\phi} E^*|_{\Sigma_\phi} \rightarrow 0.$$

The **fiber-hessian** $\mathcal{H}\phi$ of ϕ at $p \in \Sigma_\phi$ is defined as the composition $\nabla d_\theta\phi \circ \iota: E_p \rightarrow E_p^*$.

The bundle E^\perp may be identified with the pull-back $p_M^*T^*M$, giving rise to a natural projection $E^\perp \xrightarrow{p} T^*M$. On the fiber-critical set Σ_ϕ , the differential $d\phi$ defines a section of E^\perp whose composition with p we denote by $\lambda_\phi: \Sigma_\phi \rightarrow T^*M$.

Theorem 4.15 *If ϕ is nondegenerate, then the map $\lambda_\phi: \Sigma_\phi \rightarrow T^*M$ is an exact lagrangian immersion.*

Proof. Since E^\perp is the union of the conormal bundles of the fibers of p_M , it follows from Example 3.19 that E^\perp is a coisotropic submanifold of T^*B . Moreover, the characteristic distribution \mathcal{C}^\perp of E^\perp is tangent to the fibers of the mapping $E^\perp \rightarrow T^*M$.

The nondegeneracy assumption on ϕ is equivalent to the requirement that the lagrangian submanifold $(d\phi)(B)$ of T^*B be transverse to E^\perp . By Lemma 3.6(2), this implies that the section $(d\phi)(B)|_{\Sigma_\phi} = (d\phi)(B) \cap E^\perp$ is nowhere tangent to the distribution \mathcal{C}^\perp and therefore immerses into T^*M . To complete the proof, we note that the equality

$$\lambda_\phi^*\alpha_M = d\phi|_{\Sigma_\phi}$$

implies that L_ϕ is exact lagrangian. □

⁷Some readers may find it instructive to follow the ensuing discussion by writing everything in local coordinates.

From the preceding definitions, it follows that the nullity k of the fiber-hessian at $p \in \Sigma_\phi$ equals $\dim(T_p \Sigma_\phi \cap E_p)$, which in turn equals the dimension of the kernel of $(\pi \circ \lambda_\phi)_*$ on $T_p \Sigma_\phi$. This means in particular that the dimension of the fibers of B must be at least k ; if the fiber dimension of B equals k , then the phase function ϕ is said to be **reduced** at p . This occurs when its fiber-hessian vanishes at p .

Example 4.16 Suppose that $M = \mathbb{R}$ and consider the function $\phi(x, \theta) = x\theta + \theta^3/3$ on $B = \mathbb{R} \times \mathbb{R}$. Here, $d_\theta \phi = \theta^2 + x$, and $\nabla(d_\theta \phi) = 2\theta d\theta + dx$. Evidently the phase function ϕ is nondegenerate, and its fiber critical set consists of the parabola $x = -\theta^2$, which fails to be x -projectable precisely when $x = 0$. This is just the value of x for which $\phi(x, \cdot)$ acquires a degenerate critical point.

△

A phase function which generates a neighborhood of a point in a lagrangian submanifold is not unique *per se*, but it is unique up to “stable equivalence,” a concept which we now define. To begin, we introduce the following terminology.

Definition 4.17 A triple (B, p_B, ϕ) is called a **Morse family** over a manifold M if $p_B : B \rightarrow M$ is a smooth (possibly non-surjective) submersion, and ϕ is a nondegenerate phase function on B such that λ_ϕ is an embedding. A Morse family is said to be **reduced** at $p \in B$ if ϕ is a reduced phase function at p .

We will say that the lagrangian submanifold $\text{im}(\lambda_\phi) = L_\phi$ is *generated* by the Morse family (B, p_B, ϕ) . If $\iota : L \rightarrow T^*M$ is a lagrangian immersion and $p \in L$, then we denote by $\mathfrak{M}(L, \iota, p)$ the class of Morse families (B, p_B, ϕ) which generate $\iota(U)$ for some neighborhood $U \subset L$ of p . For such Morse families, we denote by $g_\phi : U \rightarrow \Sigma_\phi$ the diffeomorphism defined by $g_\phi = \lambda_\phi^{-1} \circ \iota$.

If $(B, p_B, \phi) \in \mathfrak{M}(L, \iota, p)$, then the following operations produce further elements of $\mathfrak{M}(L, \iota, p)$.

1. *Addition:* For any $c \in \mathbb{R}$, $(B, p_B, \phi + c) \in \mathfrak{M}(L, \iota, p)$.
2. *Composition:* If $p_{B'} : B' \rightarrow M$ is a second submersion and $g : B' \rightarrow B$ is a fiber-preserving diffeomorphism, then $(B', p_{B'}, \phi \circ g) \in \mathfrak{M}(L, \iota, p)$.
3. *Suspension:* The suspension of (B, p_B, ϕ) by a nondegenerate quadratic form Q on \mathbb{R}^n is defined as the Morse family comprised of the submersion $\tilde{p}_B : B \times \mathbb{R}^n \rightarrow M$ given by composing p_B with the projection along \mathbb{R}^n , together with the phase function $\tilde{\phi} = \phi + Q$. Evidently the fiber-critical set of $\tilde{\phi}$ equals the product $\Sigma_\phi \times \{0\}$, and $\lambda_{\tilde{\phi}}(b, 0) = \lambda_\phi(b)$ for all $(b, 0) \in \Sigma_{\tilde{\phi}}$. Thus $(B \times \mathbb{R}^n, \tilde{p}_B, \tilde{\phi}) \in \mathfrak{M}(L, \iota, p)$.
4. *Restriction:* If B' is any open subset of B containing $\lambda_\phi^{-1}(p)$, then the restrictions of p_B and ϕ to B' define a Morse family on M which belongs to $\mathfrak{M}(L, \iota, p)$.

These operations generate an equivalence relation among Morse families called **stable equivalence**. The central result of this section is the following.

Theorem 4.18 *Let $\iota: L \rightarrow T^*M$ be a lagrangian immersion, and let $p \in L$. Then:*

1. *The class $\mathfrak{M}(L, \iota, p)$ contains a reduced Morse family over M .*
2. *Any two members of $\mathfrak{M}(L, \iota, p)$ are stably equivalent.*

The next two subsections are devoted to the proof of this theorem.

Symplectic normal forms

The purpose of this section is to develop several results needed to prove Theorem 4.18, all of which more or less rely on the so-called *deformation method*. This method was introduced by Moser in [45] (it probably has a longer history) and applied to a variety of problems of symplectic geometry in [62]. We will prove the following theorem.

Theorem 4.19 *Let (P, ω) be a symplectic manifold, $\iota: L \rightarrow P$ a lagrangian immersion, and E a lagrangian subbundle of ι^*TP which is complementary to the image of $\iota_*: TL \rightarrow \iota^*TP$. Then there is a symplectic immersion ψ of a neighborhood U of the zero section $Z \subset T^*L$ into P such that $\psi = \iota \circ \pi$ on Z and which maps the vertical subbundle $V_Z L \subset T_Z(T^*L)$ onto E .*

As a result of the existence of compatible complex structures on ι^*P , lagrangian subbundles of ι^*P complementary to the image of $\iota_*: TL \rightarrow \iota^*TP$ always exist. Thus, Theorem 4.19 implies the following nonlinear “normal form” result, which states that, near any of its lagrangian submanifolds L , a symplectic manifold looks like a neighborhood of the zero section in T^*L .

Corollary 4.20 *If L is a lagrangian submanifold of P , then a neighborhood of L in P is symplectomorphic to a neighborhood of the zero section in T^*L , by a map which is the identity on L .*

Theorem 4.19 may be combined with Proposition 3.24 to yield:

Corollary 4.21 *The map ψ gives a 1-1 correspondence between a neighborhood of L in the space of all lagrangian submanifolds of (P, ω) and a neighborhood of the zero section in the space $Z^1(L)$ of closed 1-forms on L .*

The proof of Theorem 4.19 will rely on the following

Relative Darboux theorem [62] . *Let N be a submanifold of a manifold P , and let ω_0, ω_1 be two symplectic forms on P which coincide on $T_N P$. Then there are neighborhoods U and V of N and a diffeomorphism $f: U \rightarrow V$ such that*

1. $f^* \omega_1 = \omega_0$
2. $f|_N = \text{id}$ and $Tf|_{T_N P} = \text{id}$.

Proof. Define a 1-parameter family of closed forms on P by

$$\omega_t = \omega_0 + t(\omega_1 - \omega_0).$$

Since all ω_t agree on the submanifold N , there is a neighborhood of N in P (which for our purposes we may assume to be P itself) on which all ω_t are nondegenerate.

To find f satisfying $f^*\omega_1 = \omega_0$, we will construct a time-dependent vector field X_t for which the isotopy f_t that it generates satisfies $f_t^*\omega_t = \omega_0$ for all $t \in [0, 1]$. By the usual properties of the Lie derivative, it is necessary and sufficient that such a vector field solve the equation

$$0 = \frac{d}{dt}(f_t^*\omega_t) = f_t^* \left(\frac{d\omega_t}{dt} \right) + f_t^*(\mathcal{L}_{X_t}\omega_t) = f_t^*(\omega_1 - \omega_0 + d(X_t \lrcorner \omega_t)).$$

In order to fix N , we also want $X_t|_N = 0$. These conditions will be satisfied if we set $X_t = -\tilde{\omega}_t^{-1}(\varphi)$ for a 1-form φ on a neighborhood of N in P , vanishing on N , such that $d\varphi = \omega_1 - \omega_0$. If the submanifold N consisted of a single point, the form φ would be easy to find, since the closed 2-form $\omega_1 - \omega_0$ is locally exact by the classical Poincaré lemma. For more general submanifolds, we use a more general version of this lemma:

Relative Poincaré lemma . *Let N be a submanifold of P , and let β be a closed k -form on P which vanishes on TN . Then there is a form φ on a neighborhood of N such that $d\varphi = \beta$ and which vanishes on $T_N P$. Furthermore, if β vanishes on $T_N P$, then φ can be chosen so that its first derivatives vanish along N .*

Proof. Since the statement is local around N , we will identify P with a tubular neighborhood of N in P and let $h_t: P \rightarrow P$ be a smooth isotopy such that

$$h_1 = \text{id} \quad h_0 = \text{fiberwise projection of } P \text{ onto } N.$$

Since β vanishes on N , we have

$$\beta = h_1^*\beta - h_0^*\beta = \int_0^1 \frac{d}{dt}(h_t^*\beta) dt = \int_0^1 h_t^*(\mathcal{L}_{Y_t}\beta) dt,$$

where Y_t is the (time-dependent) vector field which generates the isotopy h_t for $t > 0$. By Cartan's formula and the fact that β is closed, the last expression is equal to

$$\int_0^1 h_t^* d(Y_t \lrcorner \beta) dt = d \left(\int_0^1 h_t^*(Y_t \lrcorner \beta) dt \right).$$

Our assertion follows if we take $\varphi = \int_0^1 h_t^*(Y_t \lrcorner \beta) dt$.

□

If (P, ω) is a symplectic manifold and $g_t: P \rightarrow P$ a continuous family of diffeomorphisms, then homotopy invariance for deRham cohomology implies that the forms $\omega_t = g_t^*\omega$ lie in the same cohomology class for all t . By the same method of proof as in the relative Darboux theorem, we can deduce the following converse result:

Corollary 4.22 (Moser [45]) *If $\{\omega_t\}$ is a family of symplectic structures on a compact manifold P and $\omega_{t_1} - \omega_{t_2}$ is exact for all t_1, t_2 , then there is a diffeomorphism $f: P \rightarrow P$ with $f_1^* \omega_1 = \omega_0$.*

By restricting to a neighborhood of a point in N , Givental [6] proves an equivalence theorem like the relative Darboux theorem under the weaker hypothesis that ω_0 and ω_1 coincide on TN , i.e. just for vectors tangent to N . In contrast to Corollary 4.22, McDuff [40] describes a family ω_t of symplectic forms on a compact six-dimensional manifold P such that $\omega_1 - \omega_0$ is exact, but there is no diffeomorphism $f: P \rightarrow P$ at all satisfying $f^* \omega_1 = \omega_0$. Of course, for intermediate t , the form $\omega_t - \omega_0$ is not exact.

Proof of Theorem 4.19. Beginning with the lagrangian immersion $\iota: L \rightarrow P$, we let $f: Z \rightarrow P$ denote the composition $\iota \circ \pi$, where Z is the zero section of T^*L and $\pi: T^*L \rightarrow L$ is the natural projection. Using Theorem 3.20, we then choose a ω -compatible complex structure J on f^*TP which satisfies $J(f_*TZ) = E$, where E is the lagrangian subbundle given in the statement of the theorem. Similarly, we choose a ω_L -compatible complex structure J' on $T_Z(T^*L)$ which rotates TZ into V_ZL . Now consider the symplectic bundle map

$$T_Z(T^*L) \xrightarrow{\tilde{f}} f^*TP$$

given by $\tilde{f}(v \oplus J'w) = f_*v \oplus J(f_*w)$ for $v, w \in TZ$. In a neighborhood $U \subset T^*L$ of Z , we can extend f to obtain an immersion $F: U \rightarrow P$ which satisfies $F_* = \tilde{f}$ on $T_Z(T^*L)$. To finish the proof of the normal form theorem, we apply the relative Darboux theorem to the forms ω_L and $F^*(\omega)$ on U .

□

Next, we turn to two generalizations of the classical Morse lemma (see Appendix B) which require the following version of Taylor's theorem.

Lemma 4.23 *Let $N \subset B$ be a submanifold defined by the vanishing of functions g_1, \dots, g_k whose differentials are linearly independent along N . If f is any function such that f and df vanish at all points of N , then there exist functions c_{ij} such that*

$$f = \sum_{i,j} c_{ij} g_i g_j$$

on a neighborhood of N .

Parametrized Morse lemma . *Let $f: M \times \mathbb{R}^k \rightarrow \mathbb{R}$ satisfy the condition that for each $x \in M$, $(x, 0)$ is a nondegenerate critical point for $f|_{\{x\} \times \mathbb{R}^k}$. Then for each $x_0 \in M$, there exists a neighborhood U of $(x_0, 0)$, a nondegenerate quadratic form Q on \mathbb{R}^k , and a diffeomorphism u of U fixing $M \times \{0\}$ and preserving fibers of the projection to M such that*

$$f(u(x, \theta)) = f(x, 0) + Q(\theta).$$

Proof. By a preliminary change of coordinates linear on fibers, we may assume that

$$f(x, \theta) = f(x, 0) + Q(\theta) + a(x, \theta)$$

where the error term $a(x, \theta)$ is $O(|\theta|^3)$ for each x . To find the diffeomorphism u we apply the deformation method. Define $f_t = f - (1 - t)a$. We seek a vector field X_t tangent to the fibers $x = \text{constant}$ which generates an isotopy u_t fixing $M \times \{0\}$ and satisfying $f_t \circ u_t = f_0$ for all $t \in [0, 1]$. This means that X_t should be of the form

$$X_t = \sum_i h_t^i \frac{\partial}{\partial \theta_i},$$

for certain smooth functions h_t^i which vanish on $M \times \{0\}$, and it should satisfy

$$0 = \frac{d}{dt}(f_t \circ u_t) = u_t^*(X_t \cdot f_t + a)$$

for all $t \in [0, 1]$. Evidently, the latter condition will be met if X_t is chosen so that

$$X_t \cdot f_t + a = 0.$$

To determine X_t , we invoke Lemma 4.23 in order to find smooth functions c_t^{ij} which vanish on $M \times \{0\}$ and satisfy

$$a = \sum c_t^{ij} \frac{\partial f_t}{\partial \theta_i} \frac{\partial f_t}{\partial \theta_j}.$$

The required condition

$$\sum_i h_t^i \frac{\partial f_t}{\partial \theta_i} = - \sum_{i,j} c_t^{ij} \frac{\partial f_t}{\partial \theta_i} \frac{\partial f_t}{\partial \theta_j}$$

will be satisfied if we set

$$h_t^i = - \sum_j c_t^{ij} \frac{\partial f_t}{\partial \theta_j},$$

which vanishes (to second order) on $M \times \{0\}$.

□

If ϕ is a phase function on $M \times \mathbb{R}^{n+k}$ whose fiber-hessian $\mathcal{H}\phi$ has rank k at a point $b \in \Sigma_\phi$, then by the lower-semicontinuity of rank $\mathcal{H}\phi$, there exists an integrable subbundle $F \xrightarrow{j} E = \ker(p_{M^*})$ such that the fiber-hessian $j^* \circ \nabla d_\theta \phi \circ j : F \rightarrow F^*$ is nondegenerate at each zero of the section $j^* \circ d_\theta \phi$ in a neighborhood of b . A change of coordinates near b and an application of the parametrized Morse lemma then proves:

Thom splitting theorem . *Let ϕ be a function on $B = M \times \mathbb{R}^{n+k}$ whose fiber-hessian $\mathcal{H}\phi$ has rank k at a point $b \in \Sigma_\phi$. Then there exists a fiber-preserving diffeomorphism g of B , a function $\eta : M \times \mathbb{R}^n \rightarrow \mathbb{R}$, and a nondegenerate quadratic form Q on \mathbb{R}^k such that*

$$(\phi \circ g)(x, \theta, \theta') = \eta(x, \theta) + Q(\theta')$$

in a neighborhood of b in B . The function η is totally degenerate in the sense that its fiber-hessian is identically zero at b .

□

Existence and stable equivalence of Morse families

Returning to the proof of Theorem 4.18, we first note the following reinterpretation of the Thom splitting theorem in the terminology of Morse families.

Theorem 4.24 *Any Morse family in $\mathfrak{M}(L, \iota, p)$ is stably equivalent to a reduced Morse family in $\mathfrak{M}(L, \iota, p)$.*

To show that $\mathfrak{M}(L, \iota, p)$ is nonempty, we first prove

Theorem 4.25 *Suppose that $\iota : L \rightarrow T^*M$ is an exact lagrangian immersion such that $\iota^*T(T^*M)$ admits a lagrangian subbundle F transverse to both ι^*VM and ι_*TL . Then L is generated by a Morse family over M .*

Proof. By Theorem 4.19, there exists a symplectic immersion ψ from a neighborhood U of the zero section $Z \subset T^*L$ into T^*M which satisfies $\psi = \iota \circ \pi$ on Z and which maps $V_Z L$ onto the subbundle F . Since ψ is symplectic, the 1-form $\psi^*\alpha_M - \alpha_L$ is closed on U . Its restriction to Z is exact because α_L is zero on Z . By the relative Poincaré lemma, there exists a function ϕ on U satisfying $d\phi = \psi^*\alpha_M - \alpha_L$.

To complete the proof, we first note that the restriction of $d\phi$ to a fiber of the composition of ψ with the projection $\pi' : T^*M \rightarrow M$ equals $-\alpha_L$. Combined with the definition of α_L , a computation (using the transversality hypothesis) now shows that there exists a neighborhood B of Z within U such that $(B, \pi'(\psi(B)), \pi' \circ \psi, \phi)$ is a Morse family such that $\Sigma_\phi = Z$ and $\lambda_\phi = \iota \circ \pi$ for the projection $\pi : T^*L \rightarrow L$.

□

Proof of Theorem 4.18. Using Lemma 3.9 we can construct a subbundle F which satisfies the hypotheses of the preceding theorem over a neighborhood of any p in L . Consequently, there exists a Morse family generating a neighborhood of p , and so the class $\mathfrak{M}(L, \iota, p)$ is nonempty. Combined with Theorem 4.24, this proves part (1) of Theorem 4.18.

To prove part (2), it suffices by Theorem 4.24 to show that any two reduced Morse families in $\mathfrak{M}(L, \iota, p)$ are stably equivalent. Since this property is local near p , we may assume that $M = \mathbb{R}^m$ and that $\phi, \tilde{\phi}$ are both defined on $\mathbb{R}^m \times \mathbb{R}^k$. Consider the mappings $F, \tilde{F} : \mathbb{R}^m \times \mathbb{R}^k \rightarrow \mathbb{R}^m \times \mathbb{R}^m$ defined by

$$F(x, y) = \left(x, \frac{\partial \phi}{\partial x}(x, y) \right) \quad \tilde{F}(x, y) = \left(x, \frac{\partial \tilde{\phi}}{\partial x}(x, y) \right).$$

By the assumption that $\phi, \tilde{\phi}$ are nondegenerate phase functions, it follows that F and \tilde{F} are embeddings near $q = \lambda_\phi^{-1}(p)$ and $\tilde{q} = \lambda_{\tilde{\phi}}^{-1}(p)$ respectively. Moreover, F coincides with the map λ_ϕ on the fiber-critical set Σ_ϕ near q , and similarly for \tilde{F} . Since the Morse families are reduced, it follows that the images of DF_q and $D\tilde{F}_{\tilde{q}}$ coincide in $T_p(\mathbb{R}^m \times \mathbb{R}^m)$. By an application of the implicit function theorem, we can therefore find a fiber-preserving map g

of $\mathbb{R}^m \times \mathbb{R}^m$ whose restriction to some neighborhood B of q is a diffeomorphism which sends Σ_ϕ to $\Sigma_{\tilde{\phi}}$ and in particular $g(q) = \tilde{q}$.

Recall from the proof of Theorem 4.15 that the identity

$$d\phi = \lambda_\phi^* \alpha_M$$

holds on the fiber critical set Σ_ϕ . Since $\lambda_\phi = \lambda_{\tilde{\phi}} \circ g$, this identity implies that $d\phi = d(\tilde{\phi} \circ g)$ on Σ_ϕ , and so the phase functions ϕ and $\tilde{\phi} \circ g$ differ by an additive constant (which we may assume to be zero) at points of Σ_ϕ . In order for g to define an equivalence between ϕ and $\tilde{\phi}$, this property would have to be valid on all of B . While there is no reason to expect this of g itself, we will construct a similar diffeomorphism with this property by appealing once again to the deformation method:

Let $\phi_0 = \phi$, $\phi_1 = \tilde{\phi} \circ g$. Then $\Sigma_{\phi_0} = \Sigma_{\phi_1} = \Sigma$ and $\lambda_{\phi_0} = \lambda_{\phi_1}$. Moreover $\phi_1 - \phi_0$ vanishes to second order along Σ , and so there exist functions c_{ij} defined near b such that

$$\phi_1 - \phi_0 = \sum_{i,j} c_{ij} \frac{\partial \phi_0}{\partial \theta_i} \frac{\partial \phi_0}{\partial \theta_j}$$

in a neighborhood of Σ_ϕ .

As before, we seek a vector field X_t generating an isotopy f_t that satisfies $f_t^* \phi_t = \phi_0$ for all $t \in [0, 1]$; to insure that each f_t fixes Σ and preserves fibers, we must also require that X_t be of the form

$$X_t = \sum_{i,j} h_{ij} \frac{\partial \phi_0}{\partial \theta_i} \frac{\partial}{\partial \theta_j}$$

for certain functions h_{ij} which vanish on Σ .

To arrive at an equation for these coefficients, we note that the equation

$$0 = \frac{d}{dt}(f_t^* \phi_t) = f_t^*(X_t \cdot \phi_t + \phi_1 - \phi_0)$$

will be satisfied provided that

$$0 = X_t \cdot \phi_t + \phi_1 - \phi_0,$$

i.e.

$$0 = \sum_{i,j} c_{ij} \frac{\partial \phi_0}{\partial \theta_i} \frac{\partial \phi_0}{\partial \theta_j} + \sum_{i,j} h_{ij} \frac{\partial \phi_0}{\partial \theta_i} \frac{\partial}{\partial \theta_j} \left(\phi_0 + t \sum_{u,v} c_{uv} \frac{\partial \phi_0}{\partial \theta_u} \frac{\partial \phi_0}{\partial \theta_v} \right).$$

This equation holds if

$$0 = C + H(I + S),$$

where $C = (c_{ij})$, $H = (h_{ij})$, and S is a matrix function which vanishes at b for all t , since ϕ_0 is reduced. Hence we can solve for H in a neighborhood of b . This completes the proof.

□

Maslov objects

If $\iota: L \rightarrow T^*M$ is any lagrangian immersion, then the symplectic vector bundle $\iota^*T(T^*M)$ over L has two lagrangian subbundles, $L_1 = \iota_*TL$ and $L_2 = \iota^*VM$. The **Maslov class** of (L, ι) is defined as the degree-1 cohomology class

$$\mu_{L, \iota} = \mu(L_1, L_2) \in H^1(L; \mathbb{R})$$

described in Example 3.21. From this definition, it follows that, unlike the Liouville classes, the Maslov classes of two immersions (L, ι) and (L, ι') are equal whenever ι and ι' are homotopic through lagrangian immersions.

Associated to any Morse family (B, p_B, ϕ) over a manifold M is an index function $\text{ind}_\phi: L_\phi \rightarrow \mathbb{Z}$ defined by

$$\text{ind}_\phi(p) = \text{index}(\mathcal{H}\phi_{\lambda_\phi^{-1}(p)}),$$

where the index of a quadratic form is, as usual, the dimension of the largest subspace on which it is negative-definite. Since the fiber-hessian is nondegenerate where L_ϕ is projectable (see the discussion following Theorem 4.15), the index function ind_ϕ is constant on any connected projectable subset of L_ϕ . From Theorem 4.18, it follows furthermore that any two index functions $\text{ind}_\phi, \text{ind}_{\phi'}$ differ by an integer on each connected component of $L_\phi \cap L_{\phi'}$.

Example 4.26 Consider $\phi(x, \theta) = \theta^3/3 + \theta(x^2 - 1)$. We have $\partial\phi/\partial\theta = \theta^2 + x^2 - 1$; thus the fiber critical set Σ_ϕ is the circle $\theta^2 + x^2 = 1$, and its image under λ_ϕ is the figure-eight. The caustic set of the projection $L \rightarrow \mathbb{R}$ consists of the points $(-1, 0), (1, 0)$. To compute the index function corresponding to ϕ , we note that

$$\frac{\partial^2\phi}{\partial\theta^2} = 2\theta;$$

consequently $\text{ind}_\phi(x, p) = 0$ for (x, p) lying on the upper right part of the curve and $\text{ind}_\phi(x, p) = 1$ for (x, p) on the upper left. Since this loop is generated by a single phase function, there is a corresponding global index function, and the index of the loop is necessarily zero.

A different situation occurs in the case of a circle. The circle has two caustics, and the jump experienced by any index function as one passes through a caustic is given by the example just computed: passing through the right caustic in the $+p$ direction decreases the index by 1, while passing through the left caustic in the same direction increases the index by 1. Traversing the circle in a counterclockwise direction, we see that the total index must change by -2 . Consequently, the circle does not admit a global generating function.

Of course, the Liouville class of the circle is also nonzero. However, it is easy to deform the circle to a closed curve with two transverse self-intersections which has zero Liouville class, but still admits no global phase function, since its Maslov class is nonzero.

△

As a degree-1 cohomology class on L , the Maslov class determines via the exponential map $\mathbb{R} \rightarrow U(1)$ an isomorphism class of flat hermitian line bundles over L . A canonical representative of this class can be constructed as follows. First, we consider the union

$$\mathfrak{M}(L, \iota) = \bigcup_{p \in L} \mathfrak{M}(L, \iota, p)$$

with the discrete topology. On the subset of $L \times \mathfrak{M}(L, \iota) \times \mathbb{Z}$ consisting of all $(p, (B, p_B, \phi), n)$ such that $(B, p_B, \phi) \in \mathfrak{M}(L, \iota, p)$, we now introduce an equivalence relation \sim by setting $(p, (B, p_B, \phi), n) \sim (\tilde{p}, (\tilde{B}, p_{\tilde{B}}, \tilde{\phi}), \tilde{n})$ provided that $p = \tilde{p}$ and

$$n + \text{ind}_\phi(p) = \tilde{n} + \text{ind}_{\tilde{\phi}}(p).$$

The quotient space with respect to this relation is a principal \mathbb{Z} -bundle $M_{L, \iota}$ over L which we will call the **Maslov principal bundle**.

Associated to the Maslov principal bundle via the representation $n \mapsto e^{i\pi n/2}$ of \mathbb{Z} in $U(1)$ is a complex line bundle $\mathcal{M}_{L, \iota}$ over L called the *Maslov line bundle*. Having discrete structure group, this line bundle carries a natural flat connection with holonomy in $\{e^{i\pi n/2}\} \simeq \mathbb{Z}_4$. Our main use of the Maslov line bundle will be to modify the half-densities on L in order to incorporate the Maslov correction into our quantization scheme.

4.4 WKB quantization

In this section we will combine the tools assembled in the preceding sections into a technique for quantizing half-densities on lagrangian submanifolds of arbitrary cotangent bundles.

The **phase bundle** associated to an immersed lagrangian submanifold $\iota: L \rightarrow T^*M$ and $\hbar > 0$ is defined as the tensor product

$$\Phi_{L, \iota, \hbar} \stackrel{\text{def}}{=} \mathcal{M}_{L, \iota} \otimes \iota^* \mathcal{E}_{M, \hbar},$$

where we recall that $\mathcal{E}_{M, \hbar}$ is the prequantum line bundle over T^*M (see Section 4.1). Observe that the product of the natural flat connections on $\mathcal{M}_{L, \iota}$ and $\iota^* \mathcal{E}_{M, \hbar}$ defines a flat connection on the phase bundle whose holonomy is represented by the mod- \mathbb{Z}_{\hbar} reduction of the real cohomology class

$$\lambda_{L, \iota} + \pi \hbar \mu_{L, \iota} / 2 \in \check{H}^1(L; \mathbb{R}),$$

which we call the **phase class** of (L, ι) . The phase bundle $\Phi_{L, \iota, \hbar}$ can be described explicitly as the collection of all quintuples $(p, t, (B, p_B, \phi), n, z)$ where $(p, t, n, z) \in L \times \mathbb{T}_{\hbar} \times \mathbb{Z} \times \mathbb{C}$ and $(B, p_B, \phi) \in \mathfrak{M}(L, \iota, p)$, modulo the equivalence relation \sim given by

$$(p, t, (B, p_B, \phi), n, z) \sim (\tilde{p}, \tilde{t}, (\tilde{B}, p_{\tilde{B}}, \tilde{\phi}), \tilde{n}, \tilde{z})$$

whenever $p = \tilde{p}$ and

$$z \cdot e^{-it/\hbar} e^{i\pi(n + \text{ind}_\phi(p))/2} = \tilde{z} \cdot e^{-i\tilde{t}/\hbar} e^{i\pi(\tilde{n} + \text{ind}_{\tilde{\phi}}(p))/2}.$$

A Morse family (B, p_B, ϕ) which generates an open subset L_ϕ of L defines a nonvanishing parallel section of $\Phi_{L, \iota, \hbar}$ over L_ϕ by

$$s_{\phi, \hbar}(p) = [p, 0, (B, p_B, \phi), 0, e^{-i\phi(y)/\hbar}],$$

where $\lambda_\phi(y) = p$. A check of these definitions shows that whenever $\iota(p) = \lambda_\phi(y) = \lambda_{\tilde{\phi}}(\tilde{y})$,

$$s_{\phi, \hbar}(p) e^{i\phi(y)/\hbar} e^{-i\pi \text{ind}_\phi(p)/2} = s_{\tilde{\phi}, \hbar}(\tilde{p}) e^{i\tilde{\phi}(\tilde{y})/\hbar} e^{-i\pi \text{ind}_{\tilde{\phi}}(\tilde{p})/2}.$$

For each $\hbar \in \mathbb{R}_+$, we denote by $\Gamma_{\text{par}}(\Phi_{L, \iota, \hbar})$ the space of parallel sections of $\Phi_{L, \iota, \hbar}$. If, for a particular \hbar , the phase class of (L, ι) is \hbar -integral, then $\Gamma_{\text{par}}(\Phi_{L, \iota, \hbar})$ is a complex vector space isomorphic to \mathbb{C} . Otherwise, $\Gamma_{\text{par}}(\Phi_{L, \iota, \hbar})$ consists of a single point (the zero section of $\Phi_{L, \iota, \hbar}$). The product

$$\Gamma_{L, \iota} = \prod_{\hbar > 0} \Gamma_{\text{par}}(\Phi_{L, \iota, \hbar})$$

then has the structure of a \mathbb{C} -module. An element $s \in \Gamma_{L, \iota}$ is then a (generally discontinuous) function which assigns to each $\hbar > 0$ an element s_\hbar in $\Gamma_{\text{par}}(\Phi_{L, \iota, \hbar})$, so that the map $p \mapsto s_\hbar(p)$ defines a parallel section of $\Phi_{L, \iota, \hbar}$. The **symbol space** of (L, ι) is defined as the complex vector space

$$\mathcal{S}_{L, \iota} \stackrel{\text{def}}{=} |\Omega|^{1/2} L \otimes_{\mathbb{C}} \Gamma_{L, \iota}.$$

The **amplitude bundle** \mathcal{A}_ϕ associated to a Morse family (B, p_B, ϕ) over a smooth manifold M is defined as the complex line bundle

$$\mathcal{A}_\phi = |\Lambda|^{1/2} B \otimes |\Lambda|^{1/2} E$$

over B , where E again denotes the subbundle $\ker(p_{B*})$ of TB . An **amplitude** is a section \mathbf{a} of \mathcal{A}_ϕ . We will say that \mathbf{a} is **properly supported** provided that the restriction of $p_B: B \rightarrow M$ to $\text{Supp}(\mathbf{a})$ is a proper map. The purpose of the space of amplitudes is to define a relation between half-densities on M and symbols on the subset L_ϕ of L generated by ϕ . To describe this relation, we begin by noting that from the exact sequence of vector bundles

$$0 \rightarrow E \rightarrow TB \rightarrow p_B^* TM \rightarrow 0$$

over B , it follows that $|\Lambda|^{1/2} B$ is naturally isomorphic to $|\Lambda|^{1/2} p_B^* TM \otimes |\Lambda|^{1/2} E$. This in turn gives rise to the natural isomorphism

$$\mathcal{A}_\phi \simeq |\Lambda|^{1/2} p_B^* TM \otimes |\Lambda| E.$$

The image of an amplitude \mathbf{a} on B under this isomorphism can be written as $p_B^* |dx|^{1/2} \otimes \sigma$, where σ is a family of 1-densities on the fibers of p_B , i.e., σ_x is a density on each nonempty $p_B^{-1}(x)$. By fiber-integration we pass to a half-density on M :

$$I_\hbar(\phi, \mathbf{a})(x) = (2\pi\hbar)^{-n/2} e^{-in\pi/4} \left(\int_{p_B^{-1}(x)} e^{i\phi/\hbar} \sigma_x \right) |dx|^{1/2},$$

where $n = \dim(p_B^{-1}(x))$, setting $I_{\hbar}(\phi, \mathbf{a})(x) = 0$ if $p_B^{-1}(x) = \emptyset$. When \mathbf{a} is properly supported, we may differentiate under the integral to conclude that $I_{\hbar}(\phi, \mathbf{a})$ is a smooth half-density on M .

To pass geometrically from \mathbf{a} to a symbol on L_{ϕ} , we first recall from Section 4.3 that the nondegeneracy of ϕ gives rise to the following exact sequence of vector bundles over the fiber-critical set Σ_{ϕ}

$$0 \rightarrow T\Sigma_{\phi} \rightarrow T_{\Sigma_{\phi}}B \xrightarrow{\nabla d_{\phi}} E^*|_{\Sigma} \rightarrow 0.$$

Since $|\Lambda|^{-1/2}E$ is naturally isomorphic to $|\Lambda|^{1/2}E^*$, this sequence induces an isomorphism of the restriction of \mathcal{A}_{ϕ} to Σ_{ϕ} with $|\Lambda|^{1/2}\Sigma_{\phi}$. If the restriction of \mathbf{a} to Σ_{ϕ} corresponds to a half-density a on Σ_{ϕ} under this isomorphism, the associated symbol on L_{ϕ} is defined as

$$\mathfrak{s}_{\mathbf{a}} \stackrel{def}{=} g_{\phi}^* a \otimes s_{\phi}.$$

(Here we recall from the discussion following Definition 4.17 that g_{ϕ} is a diffeomorphism from a neighborhood of p in L onto Σ_{ϕ} defined by the composition $\lambda_{\phi}^{-1} \circ \iota$. Also, for each \hbar , the section $s_{\phi, \hbar}$ is the canonical element of $\Gamma_{par}(\Phi_{L_{\phi, \iota, \hbar}})$ defined above). When (L, ι) is projectable, the symbol $\mathfrak{s}_{\mathbf{a}}$ and the half-density $I_{\hbar}(\phi, \mathbf{a})$ are linked by the following theorem.

Theorem 4.27 *Suppose that two Morse families $(B, p_B, \phi), (\tilde{B}, p_{\tilde{B}}, \tilde{\phi})$ generate the same projectable lagrangian embedding (L, ι) , and let $\mathbf{a}, \tilde{\mathbf{a}}$ be amplitudes on B, \tilde{B} , respectively. Then $\mathfrak{s}_{\mathbf{a}} = \mathfrak{s}_{\tilde{\mathbf{a}}}$ on L if and only if*

$$|I_{\hbar}(\phi, \mathbf{a}) - I_{\hbar}(\tilde{\phi}, \tilde{\mathbf{a}})| = O(\hbar)$$

locally uniformly on V .

Proof. By Theorem 4.18, the Morse families $(B, p_B, \phi), (\tilde{B}, p_{\tilde{B}}, \tilde{\phi})$ are stably equivalent, and so there exists a diffeomorphism $g: \Sigma_{\phi} \rightarrow \Sigma_{\tilde{\phi}}$ defined as the composition $g = g_{\tilde{\phi}} \circ g_{\phi}^{-1}$. A check of the definitions shows that the symbols $\mathfrak{s}_{\mathbf{a}}$ and $\mathfrak{s}_{\tilde{\mathbf{a}}}$ are equal precisely when

$$g^* (a \cdot e^{i\phi/\hbar} e^{-i\pi \text{ind}_{\phi}/2}) = \tilde{a} \cdot e^{i\tilde{\phi}/\hbar} e^{-i\pi \text{ind}_{\tilde{\phi}}/2},$$

where a is the half-density on Σ_{ϕ} induced by the amplitude \mathbf{a} , and similarly for \tilde{a} . Since $\tilde{\phi} = \phi \circ g$, up to a constant, Lemma B.3 implies that this occurs precisely when

$$|\det_{a_x} \mathcal{H}\phi|^{1/2} e^{-i\phi/\hbar} e^{i\pi \text{ind}_{\phi}/2} = |\det_{\tilde{a}_x} \mathcal{H}\tilde{\phi}|^{1/2} e^{-i\tilde{\phi}/\hbar} e^{i\pi \text{ind}_{\tilde{\phi}}/2}. \quad (*)$$

Since (L, ι) is projectable, we can apply the principle of stationary phase to each fiber of the projection $p_B: B \rightarrow M$ to obtain

$$I_{\hbar}(\phi, \mathbf{a})(x) = \frac{e^{i\phi/\hbar} e^{-i\pi \text{ind}_{\phi}/2}}{|\det_{a_x} \mathcal{H}\phi|^{1/2}} + O(\hbar)$$

and similarly for $I_{\hbar}(\tilde{\phi}, \tilde{\mathbf{a}})(x)$. The theorem follows by comparing these expressions with (*) above.

□

Morse families provide a general means for locally quantizing symbols on an immersed lagrangian submanifold $(L, \iota) \subset T^*M$. Suppose that (B, p_B, ϕ) is a Morse family such that the phase function ϕ generates an open subset $L_\phi \subset L$, and consider a symbol \mathfrak{s} on L supported in L_ϕ . Then there exists a unique half-density a supported in L_ϕ such that $\mathfrak{s} = s_\phi \otimes a$, and $(g_\phi)_*a$ may be canonically identified with a section of the amplitude bundle of B over the fiber-critical set Σ_ϕ , which we can extend to an amplitude \mathfrak{a} on B , compactly supported in fibers. We then set

$$I_{\hbar}(L, \iota, \mathfrak{s}) = I_{\hbar}(\phi, \mathfrak{a})(x).$$

From Theorem 4.27 we draw two conclusions about this tentative definition when L_ϕ is projectable. First, we note that if $(\tilde{B}, p_{\tilde{B}}, \tilde{\phi})$ is a second Morse family which generates L_ϕ , and $\tilde{\mathfrak{a}}$ is an amplitude on \tilde{B} obtained as above, then $\mathfrak{s}_\alpha = \mathfrak{s}_{\alpha'} = \mathfrak{s}$, and so

$$|I_{\hbar}(\phi, \mathfrak{a}) - I_{\hbar}(\tilde{\phi}, \tilde{\mathfrak{a}})| = O(\hbar).$$

Thus, $I_{\hbar}(L, \iota, \mathfrak{s})$ is well-defined, up to $O(\hbar)$ terms. Furthermore, we note that Theorem 4.27 asserts that this choice of $I_{\hbar}(L, \iota, \mathfrak{s})$ coincides with pull-back.

To quantize an arbitrary symbol \mathfrak{s} on L , we first fix a locally finite cover $\{L_j\}$ of L , such that each L_j is generated by a Morse family (B_j, p_{B_j}, ϕ_j) , and choose a partition of unity $\{h_j\}$ subordinate to $\{L_j\}$. We then set

$$I_{\hbar}(L, \iota, \mathfrak{s}) = \sum_j I_{\hbar}(L, \iota, h_j \cdot \mathfrak{s}).$$

Using the remarks above, it is easy to check that up to $O(\hbar)$ terms, this definition depends only on the semi-classical state (L, ι, \mathfrak{s}) .

Of course, although the technique above enables us to quantize all symbols on L in a consistent way, the *existence* of nonzero symbols requires that the phase bundle $\Phi_{L, \iota}$ admit nontrivial parallel sections, i.e. that the phase class of (L, ι) be \hbar -integral. For this reason, we introduce the following terminology.

Definition 4.28 *An immersed lagrangian submanifold $\iota: L \rightarrow T^*M$ is called **quantizable** if, for some $\hbar \in \mathbb{R}_+$, its phase bundle $\Phi_{L, \iota}$ admits a global parallel section, or, equivalently, if its phase class is \hbar -integral. The set of \hbar for which this condition holds will be called **admissible** for L .*

This definition is known as the **Maslov quantization condition**. Note that it is a straightforward generalization of the condition derived in the preceding section.

Example 4.29 Let N be a closed submanifold of a smooth manifold M , and let U be a tubular neighborhood of N , i.e., U is the image of the normal bundle $\nu_N \subset T_N M$ under an embedding $\psi: \nu_N \rightarrow M$ satisfying $\psi = \pi$ on the zero section of ν_N , where $\pi: \nu_N \rightarrow N$ is the natural projection. Consider the Morse family (r^*N^\perp, p_N, ϕ) , where $r = \pi \circ \psi^{-1}$ is a

retraction of U onto N , p_N denotes the natural submersion $r^*N^\perp \rightarrow M$, and $\phi: r^*N^\perp \rightarrow \mathbb{R}$ is defined by

$$\phi(p) = \langle p, \psi^{-1}(p_N(p)) \rangle.$$

Since ψ is an embedding, a computation shows that the fiber critical set of ϕ is given by $\Sigma_\phi = p_N^{-1}(N) = N^\perp$, and the map $\lambda_\phi: N^\perp \rightarrow T^*M$ equals the inclusion. Thus, the conormal bundle of N is a lagrangian submanifold of T^*M which admits a *global* generating function, and therefore both the Liouville and Maslov classes of N^\perp are zero. In particular, this implies that the conormal bundle of any submanifold of M satisfies the Maslov quantization condition.

△

Quantum states as distributions

Unfortunately, the interpretation of $I_\hbar(L, \iota, \mathfrak{s})$ at regular values of π_L is not valid at caustics. Indeed, this remark is suggested by the fact that $I_\hbar(L, \iota, \mathfrak{s})$ is smooth, whereas we saw in Section 2.2 that classical solutions to the transport equation are singular at caustic points. The basic technical difficulty is that the principle of stationary phase (the basic underpinning of Theorem 4.27) no longer applies to the integral

$$\int_{p_B^{-1}(x)} e^{i\phi/\hbar} \sigma_x$$

when x is a caustic point, since the phase function ϕ has a degenerate critical point in the fiber $p_B^{-1}(x)$.

Example 4.30 The fiber critical set of the phase function $\phi(x, \theta) = \theta^3/3 + x\theta$ consists of the parabola $x = -\theta^2$, at whose points the fiber-hessian assumes the form $\partial^2\phi/\partial\theta^2 = 2x$. The origin is therefore a degenerate critical point for ϕ , and stationary phase cannot be used to estimate the integral

$$\int_{\mathbb{R}} e^{i\theta^3/3\hbar} a(0, \theta) d\theta.$$

△

A more appropriate way to interpret the expression $I_\hbar(L, \iota, \mathfrak{s})$ in the presence of caustic points is as a family of distributional half-densities on M defined as follows. For each Morse family $(B, p_B, \phi) \in \mathfrak{M}(L, \iota)$ and compactly supported amplitude \mathfrak{a} on B , we define a distributional half-density on M by

$$\langle I_\hbar(\phi, \mathfrak{a}), u \rangle = (2\pi\hbar)^{-b/2} e^{-i\pi b/4} \int_B e^{i\phi/\hbar} \mathfrak{a} \otimes p_B^* u,$$

where $b = \dim(B)$ and $u \in |\Omega|_0^{1/2}M$. The family $I_\hbar(L, \iota, \mathfrak{s})$ then consists of all distributional half-densities I_\hbar on M obtained by choosing a locally finite open cover $\{L_j\}$ of L and a

partition of unity $\{h_j\}$ subordinate to $\{L_j\}$. Then set

$$I_{\hbar} = \sum_{j=1}^k I_{\hbar}(\phi_j, \mathbf{a}_j).$$

The family $I_{\hbar}(L, \iota, \mathfrak{s})$ consists of those distributional half-densities obtained in this way using amplitudes \mathbf{a}_j such that $\mathfrak{s}_{\mathbf{a}_j} = \mathfrak{s}$ over each L_{ϕ_j} .

Although the class $I_{\hbar}(L, \iota, \mathfrak{s})$ may appear very large, a link among its members can be described as follows. We first prove

Theorem 4.31 *Suppose that two Morse families $(B, p_B, \phi), (\tilde{B}, p_{\tilde{B}}, \tilde{\phi})$ generate the same lagrangian embedding (L, ι) , and let $\mathbf{a}, \tilde{\mathbf{a}}$ be amplitudes on B, \tilde{B} , respectively. If $\psi: M \rightarrow \mathbb{R}$ is a smooth function whose differential intersects $\iota(L)$ at exactly one point $\iota(p)$ transversely, then $\mathfrak{s}_{\mathbf{a}}(p) = \mathfrak{s}_{\tilde{\mathbf{a}}}(p)$ if and only if*

$$\left| \langle I_{\hbar}(\phi, \mathbf{a}), e^{-i\psi/\hbar} u \rangle - \langle I_{\hbar}(\tilde{\phi}, \tilde{\mathbf{a}}), e^{-i\psi/\hbar} u \rangle \right| = O(\hbar).$$

Proof. Set $R = \phi - \psi \circ p_B$. Since $\psi \circ p_B$ is constant on the fibers of p_B , the function R is a phase function having the same fiber critical set as ϕ .

As in the proof of Theorem 4.27, equality of the symbols $\mathfrak{s}_{\mathbf{a}}$ and $\mathfrak{s}_{\tilde{\mathbf{a}}}$ occurs precisely when the diffeomorphism $g: \Sigma_{\phi} \rightarrow \Sigma_{\tilde{\phi}}$ satisfies

$$g^* (a_y \cdot e^{i\phi(y)/\hbar} e^{-i\pi \text{ind}_{\phi}(p)/2}) = \tilde{a}_{\tilde{y}} \cdot e^{i\tilde{\phi}(\tilde{y})/\hbar} e^{-i\pi \text{ind}_{\tilde{\phi}}(p)/2},$$

where a is the half-density on Σ_{ϕ} induced by the amplitude \mathbf{a} , and similarly for \tilde{a} . Since $\tilde{R} = R \circ g$, up to a constant, Lemma B.3 shows that this condition is equivalent to

$$|\det_a R''(y)|^{1/2} e^{-i\phi(y)/\hbar} e^{i\pi \text{ind}_{\phi}(p)/2} = |\det_{\tilde{a}} \tilde{R}''(\tilde{y})|^{1/2} e^{-i\tilde{\phi}(\tilde{y})/\hbar} e^{i\pi \text{ind}_{\tilde{\phi}}(p)/2}.$$

Since g is fiber-preserving, we have $R(y) - \tilde{R}(\tilde{y}) = \phi(y) - \tilde{\phi}(\tilde{y})$. Moreover, it is easy to check that stable equivalence of the Morse families $(B, p_B, \phi), (\tilde{B}, p_{\tilde{B}}, \tilde{\phi})$ implies that $\text{ind}_{\phi}(p) - \text{ind}_{\tilde{\phi}}(\tilde{y}) = \text{ind} R''(y) - \text{ind} \tilde{R}''(\tilde{y})$, and so the preceding equation gives

$$\frac{e^{iR(y)/\hbar} e^{-i\pi \text{ind} R''(y)/2}}{|\det_a R''(y)|^{1/2}} = \frac{e^{i\tilde{R}(\tilde{y})/\hbar} e^{-i\pi \text{ind} \tilde{R}''(\tilde{y})/2}}{|\det_{\tilde{a}} \tilde{R}''(\tilde{y})|^{1/2}}. \quad (**)$$

The critical point y of R is nondegenerate precisely when the intersection of $\iota(L)$ with $\text{im}(d\psi)$ is transverse. In this case, an application of the principle of stationary phase gives

$$\langle I_{\hbar}(\phi, \mathbf{a}), e^{-i\psi/\hbar} u \rangle = \frac{e^{iR(y)/\hbar} e^{-i\pi \text{ind} R''(y)/2}}{|\det_a R''(y)|^{1/2}} + O(\hbar)$$

and similarly for $\langle I_{\hbar}(\tilde{\phi}, \tilde{\mathbf{a}}), e^{-i\psi/\hbar} u \rangle$. Comparing this expression with $(**)$ above completes the proof. □

Now consider a semi-classical state (L, ι, \mathfrak{s}) . As before, we let $\{L_j\}$ be a locally finite cover of L such that each L_j is generated by a Morse family (B_j, p_{B_j}, ϕ_j) over M , and choose a partition of unity $\{h_j\}$ subordinate to $\{L_j\}$. We then define

$$\langle I_{\hbar}(L, \iota, \mathfrak{s}), e^{-i\psi/\hbar}u \rangle \stackrel{def}{=} \sum_j \langle I_{\hbar}(\phi_j, \mathfrak{a}_j), e^{-i\psi/\hbar}u \rangle,$$

where \mathfrak{a}_j is the amplitude on B_j obtained in the usual way from the symbol $h_j \cdot \mathfrak{s}$ on L_j . When the image of $d\psi$ is transverse to L , then, up to $O(\hbar)$, the principal part of $\langle I_{\hbar}(L, \iota, \mathfrak{s}), e^{-i\psi/\hbar} \rangle$ depends only on the principal part of \mathfrak{s} . For a thorough exposition of this topic, we refer to [21, 28, 31].

Example 4.32 Suppose that (B, p_B, ϕ) is a Morse family over a manifold M , and suppose that $S : V \rightarrow \mathbb{R}$ is a smooth function. By setting $\tilde{S} = S \circ p_B$, we obtain a new Morse family $(B, p_B, \phi - \tilde{S})$ for which $\Sigma_{\phi} = \Sigma_{\phi - \tilde{S}}$ and $\lambda_{\phi} = f_{dS} \circ \lambda_{\phi - \tilde{S}}$, where f_{dS} is the fiberwise translation map defined by dS .

Note in particular that if $dS(p_B(x)) = \lambda_{\phi}(x)$ for some $x \in \Sigma_{\phi}$, then x is a critical point of $\phi - \tilde{S}$. This critical point is nondegenerate provided that $\lambda_{\phi - \tilde{S}}$ is transverse to the zero section of T^*M near x , or, equivalently, if λ_{ϕ} is transverse to the image of dS near x .

△

By allowing $I_{\hbar}(\phi, \mathfrak{a})$ to be a distribution, we can sometimes define it even when \mathfrak{a} does not have compact support.

Equivalent semi-classical states

A further important modification of our quantization picture is based on the conceptual distinction between an “immersion” and an “immersed submanifold”. In terms of our discussion, this means the following. If $\iota : L \rightarrow T^*M$ and $\iota' : L' \rightarrow T^*M$ are lagrangian immersions, we will say that (L, ι) and (L', ι') are *equivalent* provided that there exists a diffeomorphism $f : L \rightarrow L'$ such that $\iota = \iota' \circ f$. In this way, any lagrangian immersion $\iota : L \rightarrow T^*M$ defines an equivalence class of lagrangian immersions in T^*M which, for notational simplicity we will denote by \mathbf{L} . We refer to the equivalence class L as an **immersed lagrangian submanifold** of T^*M .

A check of the definitions shows that the main objects in our quantization scheme behave nicely with respect to this notion of equivalence. On the classical level, a diffeomorphism $f : L \rightarrow L'$ as above induces an isomorphism of symbol spaces

$$\mathcal{S}_{L, \iota} \xrightarrow{f^*} \mathcal{S}_{L', \iota'}.$$

Thus, a **symbol** on the immersed lagrangian submanifold L is well-defined as an equivalence class of symbols on the members of L . Moreover, if E is a regular value of some hamiltonian function H on T^*M , then the Hamilton-Jacobi equation

$$H \circ \iota = E$$

defines a condition on the class L which we denote by $H(L) = E$. In this case, the vector fields $X_{H,\iota}$ and $X_{H,\iota'}$ induced on $(L, \iota), (L', \iota')$ by the hamiltonian vector field of H (see Example 3.13) satisfy $X_{H,\iota'} = f_* X_{H,\iota}$, and therefore the homogeneous transport equation

$$\mathcal{L}_{X_{H,\iota}} \mathfrak{s} = 0$$

is a well-defined condition on (L, \mathfrak{s}) which we denote by $\mathcal{L}_{X_H} \mathfrak{s} = 0$.

From these remarks, we are led to view the equivalence class (L, \mathfrak{s}) as a semi-classical state in T^*M . The state is stationary with respect to a classical hamiltonian H on T^*M provided that (L, \mathfrak{s}) satisfies the associated Hamilton-Jacobi and homogeneous transport equations, as described above. On the quantum level, it is easy to check that for any members $(L, \iota, \mathfrak{s}), (L', \iota', \mathfrak{s}')$ of the equivalence class (L, \mathfrak{s}) , we have

$$I_{\hbar}(L, \iota, \mathfrak{s}) = I_{\hbar}(L', \iota', \mathfrak{s}'),$$

and thus we may define the quantization of (L, \mathfrak{s}) as the (unique) distributional half-density $I_{\hbar}(L, \mathfrak{s})$ on M obtained by quantizing any member of (L, \mathfrak{s}) . If (L, \mathfrak{s}) is a stationary semi-classical state, then $I_{\hbar}(L, \mathfrak{s})$ is a semi-classical approximate solution to the time-independent Schrödinger equation on M .

Our list of classical and quantum correspondences now assumes the form:

Object	Classical version	Quantum version
basic space	T^*M	\mathfrak{H}_M
state	(L, \mathfrak{s}) as above	distributional half-density on M
time-evolution	Hamilton's equations	Schrödinger equation
generator of evolution	function H on T^*M	operator \hat{H} on \mathfrak{H}_M
stationary state	state (L, \mathfrak{s}) satisfying $H(L) = E$ and $\mathcal{L}_{X_H} \mathfrak{s} = 0$	eigenvector of \hat{H}

With these correspondences between classical and quantum mechanics in mind, we are further led to define a semi-classical state in an *arbitrary* symplectic manifold (P, ω) to be an immersed lagrangian submanifold equipped with a half-density and possibly some “phase object” corresponding to a parallel section of the phase bundle in the cotangent bundle case. Note, however, that it is unclear what the quantum states are which are approximated by these geometric objects, since there is no underlying configuration space on which the states can live. Extending some notion of quantum state to arbitrary symplectic manifolds is one of the central goals of geometric quantization.

5 The Symplectic Category

There are many symplectic manifolds which are not cotangent bundles. For instance, an important process known as *symplectic reduction* generates many examples of such manifolds, starting with cotangent bundles. We begin this section with a discussion of reduction and then turn to the classical and quantum viewpoints in the context of general symplectic manifolds, concluding with Dirac's formulation of the quantization problem.

5.1 Symplectic reduction

The technique of symplectic reduction geometrizes the process in mechanics in which first integrals of the hamiltonian are used to eliminate variables in Hamilton's equations.

Linear symplectic reduction

Degenerate skew-symmetric bilinear forms yield symplectic vector spaces in the following way.

Lemma 5.1 *A skew-symmetric bilinear form ω on a vector space Y induces a symplectic structure on Y/Y^\perp .*

Proof. First note that ω gives rise to a skew-symmetric bilinear form ω' on Y/Y^\perp by the equation

$$\omega'([x], [y]) = \omega(x, y).$$

To prove the nondegeneracy of ω' , we use the following commutative diagram, where π is the projection.

$$\begin{array}{ccc} Y^* & \xleftarrow{\pi^*} & (Y/Y^\perp)^* \\ \tilde{\omega} \uparrow & & \uparrow \tilde{\omega}' \\ Y & \xrightarrow{\pi} & Y/Y^\perp \end{array}$$

If $\tilde{\omega}'([x]) = 0$, then $\tilde{\omega}(x) = 0$, and so $x \in Y^\perp$ and hence $[x] = 0$.

□

The symplectic quotient space Y/Y^\perp described in this lemma is known as the **(linear) reduced space** associated to Y . A special case of linear reduction arises when Y is a coisotropic subspace of a symplectic vector space and ω equals the restriction of the symplectic form to Y . Lagrangian subspaces behave remarkably well with respect to this reduction.

Lemma 5.2 *Let V be a symplectic vector space and $L, C \subset V$ a lagrangian and a coisotropic subspace respectively. Then*

$$L^C = L \cap C + C^\perp$$

is a lagrangian subspace of V contained in C , and

$$L_C = (L \cap C)/(L \cap C^\perp)$$

is a lagrangian subspace of C/C^\perp .

Proof. Note that

$$(L^C)^\perp = (L + C^\perp) \cap C.$$

Since by assumption $C^\perp \subset C$, the self-orthogonality of L^C follows from the simple fact that if E, F, G are *any* three subspaces of V , then $(E + F) \cap G = E \cap G + F$ if and only if $F \subset G$.

Next, we observe that since $L_C = L^C/C^\perp$, the second assertion follows from the equality

$$L_C^\perp = (L^C)^\perp/C^\perp.$$

□

Example 5.3 If $(V, \omega), (V', \omega')$ are symplectic vector spaces and L is a lagrangian subspace of V , then then $C = V' \oplus L$ is a coisotropic subspace of the direct sum $V' \oplus \bar{V}$ (see Example 3.4), and $C^\perp = 0 \oplus L$. Consequently, the linear reduced space C/C^\perp equals (V', ω') .

Now suppose that V and V' are of equal dimension and $T: V \rightarrow V'$ is a linear symplectic map. If $\Gamma_T \subset V' \oplus \bar{V}$ is the lagrangian subspace defined by the graph of T , then the reduced lagrangian subspace $(\Gamma_T)_C$ of C/C^\perp equals $T(L) \subset V'$.

△

Observe that if C is a coisotropic subspace of (V, ω) , then $C^\perp \oplus V/C$ also carries a natural symplectic structure induced by ω . This gives rise to the following decomposition of V .

Lemma 5.4 *If (V, ω) is a symplectic vector space with a coisotropic subspace C , then there exists a linear symplectomorphism*

$$V \rightarrow C^\perp \oplus V/C \oplus C/C^\perp.$$

Proof. Let J be a ω -compatible complex structure on V , and set $A = JC^\perp$ and $B = C \cap JC$. Then A, B are orthogonal to C, C^\perp with respect to the inner-product g_J on V . The projections $V \rightarrow V/C$ and $C \rightarrow C/C^\perp$ restricted to A, B give rise to an isomorphism

$$V = C^\perp + A + B \rightarrow C^\perp \oplus V/C \oplus C/C^\perp.$$

□

Recall that if L, L' are lagrangian subspaces of a symplectic vector space V and $W \subset V$ is a lagrangian subspace transverse to both L and L' , then there exists a natural linear symplectomorphism from V to $L \oplus L^*$ which sends W onto the subspace $0 \oplus L^*$ and L' onto the graph of some self-adjoint linear map $T: L \rightarrow L^*$. Denoting by Q_T the quadratic form on L induced by T , we define

$$\text{ind}(L, L'; W) = \text{index } Q_T \quad \text{sgn}(L, L'; W) = \text{signature } Q_T$$

These quantities will be useful in our study of the Maslov bundle under reduction, and we collect some useful facts about them in the following examples.

Example 5.5 To begin, we leave to the reader the job of checking the following elementary identities for the case when the lagrangian subspaces L, L' are themselves transverse:

$$\operatorname{sgn}(L, L'; W) = -\operatorname{sgn}(L', L; W) = -\operatorname{sgn}(L, W; L').$$

△

Example 5.6 Suppose that V is a finite-dimensional vector space with a subspace $E \subset V$ and its algebraic orthogonal $E^\perp \subset V^*$. Then $C = V \oplus E^\perp$ is a coisotropic subspace of $V \oplus V^*$ with its usual symplectic structure, and $C^\perp = E \oplus 0$. The reduced symplectic vector space is then

$$C/C^\perp = (V/E) \oplus E^\perp.$$

The composition of a self-adjoint linear map $A : V^* \rightarrow V$ with the projection $V \rightarrow V/E$ restricts to a self-adjoint linear map $A_E : E^\perp \rightarrow V/E$ (note that V/E identifies canonically with $(E^\perp)^*$). Evidently, the lagrangian subspace W of $V \oplus V^*$ given by the graph of A passes under reduction by C to the graph W_C of A_E . Now if we denote by L, L' the lagrangian subspaces $V \oplus 0$ and $0 \oplus V^*$, respectively, then $L_C = (V/E) \oplus 0$ and $L'_C = 0 \oplus E^\perp$. If Q^* denotes the quadratic form on V^* defined by A , we therefore have

$$\operatorname{sgn}(Q^*) = \operatorname{sgn}(L', W; L) \quad \text{and} \quad \operatorname{sgn}(Q^*|_{E^\perp}) = \operatorname{sgn}(L'_C, W_C; L_C).$$

From [21, p.130] we recall that if Q^* is nondegenerate and Q is the quadratic form it induces on V , then

$$\operatorname{sgn}(Q) = \operatorname{sgn}(Q|_E) + \operatorname{sgn}(Q^*|_{E^\perp}).$$

Combined with the preceding equations, this formula gives

$$\operatorname{sgn}(L', W; L) = \operatorname{sgn}(Q|_E) + \operatorname{sgn}(L'_C, W_C; L_C).$$

△

Presymplectic structures and reduction

By definition, a symplectic form is closed and nondegenerate. In some sense, the next best structure a manifold M may possess along these lines is a closed two-form ω of constant rank, i.e., with the dimension of the orthogonal $(T_x M)^\perp$ the same for all $x \in M$. In this case, ω is called a *presymplectic structure* on M with **characteristic subbundle** $(TM)^\perp$.

Theorem 5.7 *The characteristic subbundle of a presymplectic manifold (M, ω) is integrable.*

Proof. Leaving for the reader the verification that $(TM)^\perp$ is actually a subbundle of TM , we recall that Lie brackets and inner products are related by the formula

$$[X, Y] \lrcorner \omega = \mathcal{L}_X(Y \lrcorner \omega) - Y \lrcorner \mathcal{L}_X \omega.$$

If the vector field Y belongs to $(TM)^\perp$, then the first term on the right hand side of this equation vanishes identically. If X belongs to $(TM)^\perp$, then Cartan's formula combined with the assumption that ω is closed implies that ω is invariant under the flow of X . This means that $\mathcal{L}_X \omega = 0$, so the second term vanishes as well, and $[X, Y]$ lies in $(TM)^\perp$.

□

The foliation \mathcal{M}^\perp defined by the characteristic subbundle of M is known as the **characteristic foliation** of M . If the quotient space M/\mathcal{M}^\perp is a smooth manifold, then we say that M is **reducible**. A pointwise application of Lemma 5.2, together with the fact that $\mathcal{L}_X\omega = 0$ for characteristic X , shows that the presymplectic structure ω induces a smooth, nondegenerate 2-form $\omega_{\mathcal{M}}$ on M/\mathcal{M}^\perp . Since $d\omega = 0$ by hypothesis, and since the quotient map $M \rightarrow M/\mathcal{M}^\perp$ is a submersion, the form $\omega_{\mathcal{M}}$ is necessarily closed and therefore symplectic. The symplectic manifold $(M/\mathcal{M}^\perp, \omega_{\mathcal{M}})$ is called the **reduced manifold** of M .

For the most part, we will be interested in presymplectic manifolds which arise as coisotropic submanifolds of some symplectic manifold (P, ω) . Recall that a submanifold $C \subset P$ is called coisotropic if, for each $p \in C$, the tangent space T_pC contains its symplectic orthogonal

$$(T_pC)^\perp = \{v \in T_pP : \omega(v, w) = 0 \text{ for all } w \in T_pC\}.$$

In this case, we can view C as an abstract manifold and note that if ω' is the pull-back of the symplectic form ω by the natural inclusion $C \hookrightarrow P$, then $\ker(\omega') = (TC)^\perp$. Since $\dim(TC)^\perp = \dim(P) - \dim(C)$ is constant, the form ω' defines a presymplectic structure on P .

Example 5.8 As noted in Example 3.13, the hamiltonian flow associated to a function $H : \mathbb{R}^{2n} \rightarrow \mathbb{R}$ generates the characteristic foliation of any regular energy surface $H^{-1}(E)$. The corresponding reduced symplectic manifold (when it exists) will be the symplectic model for the space of quantum states of energy E .

If P, Q are symplectic manifolds, and L is a lagrangian submanifold of Q , then $P \times L$ is a coisotropic submanifold of $P \times Q$ whose characteristic foliation consists of leaves of the form $\{p\} \times L$ for $p \in P$. The product $P \times L$ is therefore reducible, and the reduction projection coincides with the usual cartesian projection $P \times L \rightarrow P$.

△

Our goal for the remainder of this section is to describe two operations on immersed lagrangian submanifolds defined by a reducible coisotropic submanifold; in effect, these operations will be nonlinear analogs of the mappings $L \mapsto L_C$ and $L \mapsto L^C$ described in the linear setting in Lemma 5.2. To begin, we recall that the **fiber product** of two maps $\iota : N \rightarrow V$ and $j : M \rightarrow V$ is defined as the subset

$$N \times_V M = (\iota \times j)^{-1}\Delta_V$$

of $N \times M$, where $\iota \times j : N \times M \rightarrow V \times V$ is the product map and $\Delta_V \subset V \times V$ is the diagonal. This gives rise to the following commutative diagram

$$\begin{array}{ccc} N \times_V M & \xrightarrow{r_M} & M \\ r_N \downarrow & & \downarrow j \\ N & \xrightarrow{\iota} & V \end{array}$$

where r_M, r_N denote the restrictions of the cartesian projections $N \times M \rightarrow M, N$ to the fiber product $N \times_V M$.

Definition 5.9 If N, M, V are smooth manifolds, then two smooth maps $\iota: N \rightarrow V$ and $j: M \rightarrow V$ are said to **intersect cleanly** provided that their fiber product $N \times_V M$ is a submanifold of $N \times M$ and

$$\iota_*TN \cap j_*TM = (j \circ r_M)_*T(N \times_V M)$$

as subbundles of $(j \circ r_M)^*TV$.

For brevity, we will say that a map $\iota: N \rightarrow V$ intersects a submanifold $W \subset V$ cleanly if ι and the inclusion map $W \hookrightarrow V$ intersect cleanly.

Example 5.10 1. If the product $\iota \times j$ of two smooth maps $\iota: N \rightarrow V$ and $j: M \rightarrow V$ is transverse to the diagonal Δ_V , then ι and j intersect cleanly. A particular case of this situation occurs when j is a submersion and ι is any smooth map.

2. Suppose that $\iota: N \rightarrow V$ and $j: M \rightarrow V$ are any smooth maps whose fiber product $N \times_V M$ is a smooth submanifold of $N \times M$. Then for any tangent vector (v, w) of $N \times_V M$, the vector $(\iota \times j)_*(v, w)$ is tangent to the diagonal in $V \times V$, and so $\iota_*v = j_*w$. Since $r_{M*}(v, w) = w$, it follows that $r_{M*}(v, w) = 0$ only if $w = 0$, in which case $\iota_*v = j_*w = 0$. Thus, if ι is an immersion, we can conclude that $r_M: N \times_V M \rightarrow M$ is an immersion as well.

In particular, if ι and j are cleanly intersecting immersions, then the maps r_N and r_M are themselves immersions.

3. A basic example of two maps which do not intersect cleanly is provided by two embedded curves in the plane which intersect tangentially at a single point.

△

The usefulness of clean intersection in symplectic geometry lies in its compatibility with symplectic orthogonalization, as in the following lemma.

Lemma 5.11 Let P be a symplectic manifold. If a lagrangian immersion $\iota: L \rightarrow P$ intersects a coisotropic submanifold $C \subset P$ cleanly, then the map $(p_C \circ r_C): L \times_P C \rightarrow C/C^\perp$ has constant rank.

Proof. Let F, G denote the lagrangian and coisotropic subbundles r_L^*TL and r_C^*TC of the symplectic vector bundle r_C^*TP over the fiber product $L \times_P C$. The kernel of $(p_C \circ r_C)_*$ then equals

$$G^\perp \cap r_{C*}T(L \times_P C).$$

The clean intersection hypothesis implies that $r_{C*}T(L \times_P C) = F \cap G$, and so the basic properties of symplectic orthogonalization applied fiberwise to these bundles give

$$\begin{aligned} \ker(p_C \circ r_C)_* &= G^\perp \cap G \cap F \\ &= G^\perp \cap F \\ &= (G + F)^\perp. \end{aligned}$$

Since $F \cap G$ has constant rank, the internal sum $F + G$ and its orthogonal $(F + G)^\perp$ also have constant rank, indicating that the dimension of $\ker(p_C \circ r_C)_*$ is independent of the base point in $L \times_P C$. In fact, we have $\dim(\ker(p_C \circ r_C)_*) = \dim(L) + \dim(L \times_P C) - \dim(C)$.

□

If the quotient of $L \times_P C$ by the fibers of the map $p_C \circ r_C$ is a smooth, Hausdorff manifold, then we obtain an induced immersion ι_C from the reduced space L_C of $L \times_P C$ into C/\mathcal{C}^\perp . In this case, we define L^C as the fiber product of ι_C and the quotient map $p_C: C \rightarrow C/\mathcal{C}^\perp$. By $j_C: L^C \rightarrow C/\mathcal{C}^\perp$ we denote the restriction of the cartesian projection $L_C \times C \rightarrow C$ to L^C .

Recall from Chapter 4 that an immersed submanifold N of a manifold V is an equivalence class N of immersions, where $\iota: N \rightarrow V$ is considered equivalent to $\iota': N' \rightarrow V$ provided that there exists a diffeomorphism $f: N \rightarrow N'$ which intertwines ι and ι' , i.e. $\iota = \iota' \circ f$. Evidently, if a member $\iota: N \rightarrow V$ of the equivalence class N intersects a smooth map $j: M \rightarrow V$ cleanly, then the same is true of any other member of N . In this case, we will say that the immersed submanifold N and the map $j: M \rightarrow V$ intersect cleanly.

With these remarks in mind, we will say that a coisotropic submanifold C and an immersed lagrangian submanifold L in a symplectic manifold P form a *reducible pair* (C, L) provided that C is reducible, and L, C intersect cleanly. We leave it to the reader to check that any reducible pair (C, L) in P induces immersed submanifolds L_C and L^C of C/\mathcal{C}^\perp and P , respectively.

Theorem 5.12 *If (C, L) is a reducible pair in a symplectic manifold P , then L_C and L^C are immersed lagrangian submanifolds of C/\mathcal{C}^\perp and P , respectively.*

Proof. Let $\iota: L \rightarrow P$ be a fixed element of L . Since L_C is the quotient of $L \times_P C$ by the fibers of $p_C \circ r_C$, the induced map $\iota_C: L_C \rightarrow C/\mathcal{C}^\perp$ is a smooth immersion. A pointwise application of Lemma 5.2 shows that this immersion is lagrangian.

Since $\iota_C: L_C \rightarrow C/\mathcal{C}^\perp$ is an immersion, Example 5.10(2) implies that the map $j_C: L^C \rightarrow C$ is a smooth immersion, and a pointwise application of Lemma 5.2 implies that this immersion is lagrangian.

Finally, since the immersed submanifolds L_C and L^C are well-defined by the reducible pair (C, L) , it follows that these are immersed lagrangian submanifolds.

□

Example 5.13 If (M, g) is a riemannian manifold, all of whose geodesics are closed and have the same length, then for each $E > 0$, the constant energy hypersurface $C = k_M^{-1}(E)$ (see Section 3.2) is reducible, and the reduced manifold $CM = C/\mathcal{C}^\perp$ is the space of oriented geodesics on M . The tangent space to CM at a point p identifies naturally with the space of Jacobi fields normal to the geodesic represented by p . Moreover, it is easy to check that for any pair J_1, J_2 of normal Jacobi fields along p , the symplectic form on CM is given in terms of the metric by

$$\omega(J_1, J_2) = g(J_1, J_2') - g(J_1', J_2).$$

For further details, see [12].

The conormal bundle N^\perp to a smooth submanifold $N \subset M$ is a lagrangian submanifold of T^*M which intersects C cleanly. More precisely, $N^\perp \cap C$ is a sphere bundle over N transverse to the characteristic foliation of C . From Theorem 5.12 it therefore follows that the space of geodesics in M normal to N at some point comprises an immersed lagrangian submanifold of CM .

To give a concrete case of this example, we consider the sphere S^n equipped with its standard metric. Each oriented geodesic in S^n identifies naturally with an oriented 2-dimensional subspace of \mathbb{R}^{n+1} , and therefore CS^n is represented by the grassmannian $Gr^+(2, n+1)$.

△

Example 5.14 Note that the standard symplectic structure on \mathbb{R}^{2n+2} is equivalent to the imaginary part of the standard hermitian metric on \mathbb{C}^{n+1} . The unit sphere S^{2n+1} is a coisotropic submanifold of \mathbb{C}^{n+1} , the reduction of which is complex projective space $\mathbb{C}P^n$. Since the symplectic form on \mathbb{C}^{n+1} is invariant under unitary transformations, the resulting symplectic form on $\mathbb{C}P^n$ is invariant under transformations induced by unitary transformations.

The maximal real subspace $\mathbb{R}^{n+1} \subset \mathbb{C}^{n+1}$ is a lagrangian submanifold of \mathbb{C}^{n+1} whose intersection with S^{2n+1} is clean and coincides with the real unit sphere S^n . From Theorem 5.12, it follows that the real projective space $\mathbb{R}P^n$ is a lagrangian submanifold of $\mathbb{C}P^n$.

△

Conormal submanifolds

A particular example of symplectic reduction which we will use in the next chapter begins with the following set-up. Suppose that M is a smooth manifold, and let $N \subset M$ be a smooth submanifold equipped with a tangent distribution \mathcal{F} , i.e., a subbundle of $TN \subset TM$. By C_N we denote the *annihilator* of \mathcal{F} , that is

$$C_N = \{(x, p) \in T^*M : x \in N, \mathcal{F}_x \subset \ker(p)\}.$$

By definition, C_N is a subbundle of T_N^*M , but we wish to consider various properties of C_N when it is viewed as a smooth *submanifold* of the symplectic manifold (T^*M, ω_M) . We begin with the following lemma.

Lemma 5.15 *In the notation above, the annihilator C_N is a coisotropic submanifold of T^*M if and only if the distribution \mathcal{F} is integrable.*

Proof. If \mathcal{F} is integrable, then we can consider its leaves as submanifolds of M via the inclusion $N \rightarrow M$. The conormal bundles of the leaves define a foliation of C_N by lagrangian submanifolds; by Lemma 3.6, this implies that C_N is coisotropic.

Now let Z denote the intersection of C_N with the zero section of T^*M . Then the natural projection $T^*M \rightarrow M$ maps Z diffeomorphically onto N , and moreover, $T_Z C_N = TN \oplus \mathcal{F}^\perp$,

where \mathcal{F}^\perp is the subbundle of $V_Z M$ which annihilates \mathcal{F} . From the natural lagrangian splitting of $T(T^*M)$ along the zero section of T^*M , it follows that

$$(T_Z C_N)^\perp = \mathcal{F} \oplus (TN)^\perp,$$

and thus the intersections of Z with the isotropic leaves of C_N are integral manifolds of the distribution \mathcal{F} . This completes the proof. \square

Lemma 5.16 *If the leaf space N/\mathcal{F} is a smooth, Hausdorff manifold, then the reduced space C_N/\mathcal{C}^\perp is canonically symplectomorphic to $T^*(N/\mathcal{F})$.*

Proof. Since, by definition, each $(x, p) \in C_N$ contains the subspace \mathcal{F}_x of $T_x N$ in its kernel, it follows that (x, p) defines an element of $T^*(N/\mathcal{F})$. This gives a well-defined surjective submersion $f: C_N \rightarrow T^*(N/\mathcal{F})$, and it is easy to check that the level sets of this map are the fibers of the projection $C_N \rightarrow C_N/\mathcal{C}^\perp$. Thus we get a diffeomorphism $C_N/\mathcal{C}^\perp \rightarrow T^*(N/\mathcal{F})$.

To prove that the preceding diffeomorphism is symplectic, it suffices to note that by the commutative diagram

$$\begin{array}{ccc} C_N & \xrightarrow{f} & T^*(N/\mathcal{F}) \\ \downarrow & & \downarrow \\ N & \longrightarrow & N/\mathcal{F} \end{array}$$

the Liouville forms of T^*M and $T^*(N/\mathcal{F})$ are related by $\alpha_M|_{C_N} = f^* \alpha_{N/\mathcal{F}}$. \square

Example 5.17 Suppose that \mathcal{F} is a foliation of a manifold N . By Lemma 5.15, it follows that the annihilator $C_N \subset T^*N$ of \mathcal{F} is a coisotropic submanifold which, by definition, contains the conormal bundle F^\perp to any leaf F of \mathcal{F} . Since F^\perp is a lagrangian submanifold of T^*N , Lemma 3.6 implies that F^\perp is foliated by isotropic leaves of C_N . Moreover, the proof of Lemma 5.16 shows that each such leaf projects diffeomorphically onto F , and so the bundle F^\perp is equipped in this way with a flat Ehresmann connection (see Ehresmann [22]) known as the **Bott connection** associated to the foliation \mathcal{F} . \triangle

Example 5.18 As a special case of Example 5.17, suppose that (B, p_B, ϕ) is a Morse family over a manifold M , and let $C_B \subset T^*B$ be the conormal bundle to the fibers of the submersion $p_B: B \rightarrow V$. If $L \subset T^*B$ is the lagrangian submanifold given by the image of the differential $d\phi$, then the nondegeneracy condition on ϕ implies that L intersects C_B transversally along a submanifold I which maps diffeomorphically onto the fiber critical set Σ_ϕ under the natural projection $\pi: T^*B \rightarrow B$. The lagrangian immersion $\lambda_\phi: \Sigma_\phi \rightarrow T^*M$ equals the composition

$$\Sigma_\phi \xrightarrow{\pi^{-1}} L \cap C_B \xrightarrow{p_{C_B}} C_V/\mathcal{C}^\perp.$$

In other words, Theorem 4.15 can be reformulated as an application of Theorem 5.12 when the lagrangian submanifold $\text{im}(d\phi)$ is transverse to E , i.e., when ϕ is nondegenerate.

△

We remark that since C_N is a bundle over N , the intersection $TC_N \cap VM$ is a (constant rank) subbundle of TC_N . A check of the definitions shows that this subbundle is mapped onto $V(N/\mathcal{F})$ under the quotient map $C_N \rightarrow T^*(N/\mathcal{F})$ described above.

Example 5.19 An important special case of Examples 5.6 and 5.18 arises when (B, V, p_V, ϕ) is a Morse family which generates a lagrangian submanifold in $L \subset T^*M$ and $x \in B$ is a (nondegenerate) critical point of ϕ , so that $p = \lambda_\phi(x)$ lies in the zero section Z_M of T^*M . Denoting by x' the point in the zero section of T^*B lying over x , we recall that $T_{x'}(T^*B)$ admits the natural lagrangian splitting $T_x B \oplus T_x^* B$. Let $E \subset T_x B$ be the kernel of the linearized projection $p_{B^*}: TB \rightarrow TM$.

Now suppose that Q is the nondegenerate quadratic form on $T_x B$ defined by the hessian of ϕ . By definition, the restriction of Q to E equals the form induced by the fiber-hessian $\mathcal{H}\phi_x$ of ϕ at x , and so

$$\text{ind}(Q|_E) = \text{ind } \mathcal{H}\phi_x.$$

On the other hand, it is easy to check that the tangent space to the conormal submanifold C_V defined in Example 5.18 at x equals the coisotropic subspace $C = T_x B \oplus E^\perp$, so that the last formula in Example 5.6 implies that the Morse index $\iota_M(x, \phi)$ of the critical point x satisfies

$$\iota_M(x, \phi) = \text{ind } \mathcal{H}\phi_x + \text{ind}(V_p M, T_p L; T_p Z_M).$$

△

Reduction and the Maslov bundle

Using linear reduction we can give a useful alternative description of the Maslov bundle of an immersed lagrangian submanifold L in T^*M . Given two lagrangian subspaces L, L' of a symplectic vector space V , we denote by $\mathcal{L}_{L, L'}$ the subset of the lagrangian grassmannian $\mathcal{L}(V)$ comprised of those lagrangian subspaces which are transverse to both L and L' . For $W, W' \in \mathcal{L}_{L, L'}$, the **cross index** of the quadruple $(L, L'; W, W')$ is defined as the integer

$$\sigma(L, L'; W, W') = \text{ind}(L, L'; W) - \text{ind}(L, L'; W').$$

Immediately from this definition and Example 5.5, we obtain the following ‘‘cocycle’’ property of the cross index.

Lemma 5.20 *If L, L' are lagrangian subspaces of V and $W, W', W'' \in \mathcal{L}_{L, L'}$, then*

$$\sigma(L, L'; W, W') + \sigma(L, L'; W', W'') + \sigma(L, L'; W'', W) = 0.$$

□

We denote by $\mathcal{F}_{L,L'}(V)$ the space of functions $f: \mathcal{L}_{L,L'} \rightarrow \mathbb{Z}$ such that

$$f(W) - f(W') = \sigma(L, L'; W, W')$$

for all $W, W' \in \mathcal{L}_{L,L'}$. Since any such function is determined up to an additive constant by its value at a single point of $\mathcal{L}_{L,L'}$, it follows that \mathbb{Z} acts simply and transitively by addition on $\mathcal{F}_{L,L'}(V)$.

Now we may assign a principal \mathbb{Z} bundle to a triple (E, λ, λ') , where E is a symplectic vector bundle and λ, λ' are lagrangian subbundles. Associated to E is its lagrangian grassmannian bundle $\mathcal{L}(E)$ whose fiber over $x \in M$ is simply $\mathcal{L}(E_x)$. The lagrangian subbundles λ, λ' then correspond to smooth sections of $\mathcal{L}(E)$, and we denote by $M_{\lambda, \lambda'}(E)$ the principal \mathbb{Z} bundle whose fiber over $x \in M$ equals $\mathcal{F}_{\lambda_x, \lambda'_x}(E_x)$.

A special case of this set-up occurs when L is an immersed lagrangian submanifold of T^*M so that $E = \iota^*T(T^*M)$ is a symplectic vector bundle over L . Natural lagrangian subbundles of E are then given by $\lambda = \iota_*TL$ and $\lambda' = \iota^*VM$.

Theorem 5.21 *In the notation above, the Maslov bundle $M_{L, \iota}$ is canonically isomorphic to $M_{\lambda, \lambda'}(E)$.*

Proof. Let (B, p_B, ϕ) be a Morse family over a manifold M , fix a point x in the fiber critical set Σ_ϕ , and set $p = \lambda_\phi(x)$. From Example 4.32, we recall that if $S: M \rightarrow \mathbb{R}$ is a smooth function such that the image of dS intersects $\lambda_\phi(\Sigma_\phi)$ transversely at p , then x is a nondegenerate critical point of $\phi - \tilde{S}$, where $\tilde{S} = S \circ p_B$. From Example 5.19 it follows that the Morse index $\iota_M(x, \phi - \tilde{S})$ at x is given by

$$\iota_M(x, \phi - \tilde{S}) = \text{ind } \mathcal{H}\phi_x + \text{ind}(V_{p'}M, T_{p'}L_{\phi-S}; T_{p'}Z_M),$$

where $p' = p - dS(p_V(x))$ and $L_{\phi-\tilde{S}}$ is the lagrangian submanifold generated by $\phi - \tilde{S}$. Fiberwise translation by dS does not change the index on the right, and moreover it maps $T_{p'}L_{\phi-S}$ to T_pL_ϕ , the subspace $V_{p'}M$ to the vertical subspace V_pM at p , and $T_{p'}Z_M$ to the tangent space Y of the image of dS at p . Thus, the preceding equation and Example 5.5 imply

$$\iota_M(x, S) = \text{ind } \mathcal{H}\phi_x + \text{ind}(T_pL_\phi, V_pM; Y).$$

Now if Y is any lagrangian subspace of $T_p(T^*M)$ transverse to V_pM , then there exists a function $S: V \rightarrow \mathbb{R}$ such that $Y = \text{im } dS(p_V(x))$. A function $f_\phi: \mathcal{L}_{T_pL_\phi, V_pM} \rightarrow \mathbb{Z}$ is obtained by setting

$$f_\phi(Y) = \iota_M(x, \phi - S_Y).$$

From the discussion above, it follows that this function satisfies

$$f_\phi(Y) - f_\phi(Y') = \sigma(T_pL_\phi, V_pM; Y, Y').$$

Using these remarks, we define a map $L \times \mathfrak{M}(L, \iota) \times \mathbb{Z} \rightarrow M_{\lambda, \lambda'}(E)$ by sending the triple $(p, (B, p_B, \phi), n)$ to the function (p, f_ϕ) . From the preceding paragraph and the definition of $M_{\lambda, \lambda'}(E)$, it follows that this map passes to an isomorphism of principal \mathbb{Z} bundles $M_{L, \iota} \rightarrow M_{\lambda, \lambda'}(E)$.

□

We refer to [3] and [31] for a proof that the Maslov bundle described above is equivalent to the Maslov bundle constructed using index functions in Chapter 4. For the remainder of this section, our goal is to establish the following property of the bundle $M_{\lambda,\lambda'}(E)$ under symplectic reduction. The proofs are rather technical and can be passed over in a first reading.

Theorem 5.22 *Suppose that E is a symplectic vector bundle over M with lagrangian subbundles λ, λ' and a coisotropic subbundle η . If $\lambda \cap \eta$ and $\lambda' \cap \eta$ have constant rank, then $M_{\lambda,\lambda'}(E)$ is canonically isomorphic to $M_{\lambda_\eta,\lambda'_\eta}(\eta/\eta^\perp)$.*

The starting point of the proof is the following special case of the theorem.

Lemma 5.23 *Suppose that E is a symplectic vector bundle with lagrangian subbundles λ, λ' and a coisotropic subbundle η . If $\lambda, \lambda' \subset \eta$, then $M_{\lambda,\lambda'}(E)$ is canonically isomorphic to $M_{\lambda_\eta,\lambda'_\eta}(\eta/\eta^\perp)$.*

Proof. First consider a symplectic vector space V with lagrangian subspaces L, L' and a coisotropic subspace C . If $L, L' \subset C$, then there exists a natural linear symplectomorphism of reduced spaces

$$(L_C + L'_C)/(L_C + L'_C)^\perp \rightarrow (L + L')/(L + L')^\perp$$

which maps $L_C/(L_C + L'_C)^\perp$ onto $L/(L + L')^\perp$ and $L'_C/(L_C + L'_C)^\perp$ onto $L'/(L + L')^\perp$. Moreover, if $W \in \mathcal{L}_{L,L'}$, then L_C is transverse to L_C, L'_C and $W_C/(L_C + L'_C)^\perp$ maps onto $W/(L + L')^\perp$ under the symplectomorphism above. Thus,

$$\text{ind}(L, L'; W) = \text{ind}(L_C, L'_C; W_C).$$

Using this remark, we can define the isomorphism $M_{\lambda,\lambda'}(E) \rightarrow M_{\lambda_\eta,\lambda'_\eta}(\eta/\eta^\perp)$ fiberwise by sending $(x, f) \in M_{\lambda,\lambda'}(E)$ to (x, f_η) , where

$$f_\eta(W_{\eta_x}) = f(W)$$

for each $W \in \mathcal{L}_{\lambda_x,\lambda'_x}$.

□

Lemma 5.24 *Suppose that E is a symplectic vector bundle over Y . If λ, λ' are lagrangian subbundles such that $\lambda \cap \lambda'$ has constant rank, then $M_{\lambda,\lambda'}(E)$ is canonically isomorphic to $Y \times \mathbb{Z}$.*

Proof. If λ, λ' are fiberwise transverse, then there exists a vector bundle symplectomorphism $E \rightarrow \lambda \oplus \lambda^*$ which maps λ onto $\lambda \oplus 0$ and λ' onto $0 \oplus \lambda^*$. Any choice of a positive-definite quadratic form on λ is induced by a section T of $\text{Hom}(\lambda, \lambda^*)$ comprised of self-adjoint maps.

The lagrangian subbundle λ'' of E which maps to the subbundle of $\lambda \oplus \lambda^*$ defined by the graph of T is then transverse to λ, λ' and clearly satisfies

$$\text{ind}(\lambda_x, \lambda'_x; \lambda''_x) = 0$$

for all $x \in Y$. A natural section $s(x) = (x, f)$ of $M_{\lambda, \lambda'}(E)$ is therefore given by

$$f(\lambda''_x) = 0.$$

In general, our hypotheses imply that $\eta = \lambda + \lambda'$ is a coisotropic subbundle of E which contains the lagrangian subbundles λ and λ' . By Lemma 5.23, the bundle $M_{\lambda, \lambda'}(E)$ is canonically isomorphic to $M_{\lambda_\eta, \lambda'_\eta}(\eta/\eta^\perp)$. Since $\lambda_\eta, \lambda'_\eta$ are transverse lagrangian subbundles of η/η^\perp , the latter bundle has a canonical trivialization, and the assertion follows. \square

Proof of Theorem 5.22. An application of the cocycle identity of Lemma 5.20 shows that if $\lambda, \lambda', \lambda''$ are lagrangian subbundles of a symplectic vector bundle E , then $M_{\lambda, \lambda'}(E)$ is canonically isomorphic to the principal-bundle product (see Appendix D for the definition) $M_{\lambda, \lambda''}(E) \times M_{\lambda'', \lambda'}(E)$. Applying this remark, we obtain the following canonical isomorphisms:

$$\begin{aligned} M_{\lambda, \tilde{\lambda}}(E) &\simeq M_{\lambda, \lambda^\eta}(E) \times M_{\lambda^\eta, \tilde{\lambda}}(E) \\ &\simeq M_{\lambda, \tilde{\lambda}}(E) \times (M_{\lambda^\eta, \tilde{\lambda}^\eta}(E) \times M_{\tilde{\lambda}, \tilde{\lambda}^\eta}(E)). \end{aligned}$$

From the assumption that $\lambda \cap \eta$ has constant rank, it follows that $\lambda \cap \lambda^\eta$ has constant rank, and thus $M_{\lambda, \lambda^\eta}(E)$ is canonically trivial by Lemma 5.24. Similarly, $M_{\tilde{\lambda}, \tilde{\lambda}^\eta}(E)$ is canonically trivial. Combined with the canonical isomorphisms above, this shows that $M_{\lambda, \tilde{\lambda}}(E)$ is canonically isomorphic to $M_{\lambda^\eta, \tilde{\lambda}^\eta}(E)$. Since the lagrangian subbundles $\lambda^\eta, \tilde{\lambda}^\eta$ of E are contained in the coisotropic subbundle η , Lemma 5.23 implies that $M_{\lambda^\eta, \tilde{\lambda}^\eta}(E)$ is in turn canonically isomorphic to $M_{\lambda_\eta, \tilde{\lambda}_\eta}(\eta/\eta^\perp)$, completing the proof. \square

Example 5.25 Suppose that E, E' are trivial symplectic vector bundles over M with lagrangian subbundles $\lambda \subset E$ and $\lambda' \subset E'$, and let T be a section of the bundle $\text{Hom}_\omega(E, E')$ of symplectic vector bundle maps. The graph of the section T defines a lagrangian subbundle $\lambda_{\Gamma(T)}$ of $E' \oplus \bar{E}$, while the restriction of T to the subbundle λ defines a lagrangian subbundle $T(\lambda)$ of E' .

It is easy to check that the lagrangian subbundles $\lambda_{\Gamma(T)}$ and $\lambda' \oplus \lambda$ intersect the coisotropic subbundle $E' \oplus \lambda$ along constant rank subbundles, whereas their reductions by $E' \oplus \lambda$ are given by $T(\lambda)$ and λ' , respectively. Thus, Theorem 5.22 implies that $M_{\lambda_{\Gamma(T)}, \lambda' \oplus \lambda}(E' \oplus E)$ and $M_{T(\lambda), \lambda'}(E')$ are canonically isomorphic. \triangle

Example 5.26 Any symplectomorphism $F: T^*M \rightarrow T^*N$ induces a lagrangian embedding $\iota_F: T^*M \rightarrow T^*(M \times N)$, defined as the composition of the graph $\Gamma_F: T^*M \rightarrow T^*N \times \overline{T^*M}$ of F with the Schwartz transform $S_{M,N}: T^*N \times \overline{T^*M} \rightarrow T^*(M \times N)$. A second lagrangian embedding $z_F: M \rightarrow T^*N$ is defined by the composition of F with the zero section $s_0: M \rightarrow T^*M$. Observe that these definitions imply that the Liouville class λ_{M,z_F} of the embedding (M, z_F) identifies with λ_{T^*M,ι_F} under the natural isomorphism $H^1(M; \mathbb{R}) \simeq H^1(T^*M; \mathbb{R})$ induced by the projection $T^*M \rightarrow M$.

To see that a similar relation holds for the Maslov classes of (M, z_F) and (T^*M, ι_F) , consider the symplectic vector bundles $E = s_0^* T(T^*M)$ and $E' = z_F^* T(T^*N)$ over M , along with the section T of $\text{Hom}_\omega(E, E')$ defined by the restriction of F_* to $T_Z(T^*M)$. Moreover, we have the lagrangian subbundles $\lambda = s_0^* VM$ of E and $\lambda' = z_F^* VN$ of E' , and the lagrangian subbundles $\lambda_{\Gamma(T)} \subset E' \oplus \overline{E}$ and $T(\lambda) \subset E'$, as described in Example 5.25. Theorem 5.21 shows that $s_0^* M_{T^*M,\iota_F}$ is canonically isomorphic to $M_{\lambda_{\Gamma(T)}, \lambda' \oplus \lambda}(E' \oplus E)$ and that M_{M,z_F} is canonically isomorphic to $M_{T(\lambda), \lambda'}(E')$. Thus, it follows from Example 5.25 that the Maslov classes are also related by pull-back: $s_0^* \mu_{T^*M,\iota_F} = \mu_{M,z_F}$.

△

Example 5.27 We can apply the conclusion of the preceding example to fiber-preserving symplectomorphisms $f: T^*M \rightarrow T^*N$. From Section 3.2, we recall that any such f is equal to a fiberwise translation of T^*M by a closed 1-form β on M composed with the cotangent lift of some diffeomorphism $N \rightarrow M$. Evidently, the phase class of the induced lagrangian embedding (M, z_f) equals $[\beta] \in \check{H}^1(M; \mathbb{R})$, and thus $\hbar \in \mathbb{R}_+$ is admissible for the lagrangian embeddings $(M, z_f), (T^*M, \iota_f)$ if and only if $[\beta]$ is \hbar -integral.

△

5.2 The symplectic category

To systematize the geometric aspects of quantization in arbitrary symplectic manifolds, we now introduce the symplectic category \mathfrak{S} . As objects of \mathfrak{S} we take the class of all smooth, finite-dimensional symplectic manifolds. Given two objects (P, ω) (Q, ω') of \mathfrak{S} , we define their product as the symplectic manifold $(P \times Q, \pi_1^* \omega + \pi_2^* \omega')$, where π_1, π_2 denote the cartesian projections. The symplectic dual of an object (P, ω) is the object $(P, -\omega)$.

From Lemma 3.14, we recall that a smooth diffeomorphism from a symplectic manifold P to a symplectic manifold Q is a symplectomorphism if and only if its graph is a lagrangian submanifold of $Q \times \overline{P}$. More generally, an immersed lagrangian submanifold of $Q \times \overline{P}$ is called a **canonical relation** from P to Q . The morphism set $\text{Hom}(P, Q)$ is then defined to consist of all canonical relations in $Q \times \overline{P}$. Since immersed lagrangian submanifolds of the product $Q \times \overline{P}$ coincide with those of its dual, we can therefore define the **adjoint** of a canonical relation $L \subset \text{Hom}(P, Q)$ as the element $L^* \in \text{Hom}(Q, P)$ represented by the same equivalence class L of immersions into the product $P \times Q$.

Composition of morphisms is unfortunately not defined for all $L_1 \in \text{Hom}(P, Q)$ and $L_2 \in \text{Hom}(Q, R)$, and so \mathfrak{S} is therefore not a true category. Nevertheless, we can describe

a sufficient condition for the composability of two canonical relations as follows. By Example 5.8, the product $R \times \Delta_Q \times \overline{P}$ is a reducible coisotropic submanifold of $R \times \overline{Q} \times Q \times \overline{P}$, where Δ_Q denotes the diagonal in $\overline{Q} \times Q$. The product $L_2 \times L_1$ of the canonical relations L_1 and L_2 is a well-defined immersed lagrangian submanifold of $R \times \overline{Q} \times Q \times \overline{P}$, and we will call $L_2 \times L_1$ **clean** if $(R \times \Delta_Q \times \overline{P}, L_2 \times L_1)$ is a reducible pair. Applying Theorem 5.12, we then obtain

Proposition 5.28 *If $L_2 \times L_1$ is clean, then $L_2 \circ L_1$ is an immersed lagrangian submanifold of $R \times \overline{P}$, i.e. $L_2 \circ L_1 \in \text{Hom}(P, R)$.*

□

Associativity of compositions holds in the symplectic category in the sense that for canonical relations L_1, L_2, L_3 , the equation

$$L_1 \circ (L_2 \circ L_3) = (L_1 \circ L_2) \circ L_3$$

is valid provided that both sides are defined.

Among the members of \mathfrak{S} , there is a minimal object Z , the zero-dimensional symplectic manifold consisting of a single point $*$ equipped with the null symplectic structure. Morphisms from Z to any other object $P \in \mathfrak{S}$ identify naturally with immersed lagrangian submanifolds of P . Thus, the “elements” of P are its immersed lagrangian submanifolds, and in particular, we may identify the set $\text{Hom}(P, Q)$ of morphisms with the “elements” of $Q \times \overline{P}$ for any $P, Q \in \mathfrak{S}$.⁸

A canonical relation $L \in \text{Hom}(P, Q)$ is said to be a **monomorphism** if the projection of L onto P is surjective and the projection of L onto Q is injective, in the usual sense. Equivalently, L is a monomorphism if $L^* \circ L = \text{id}_P$. Dually, one defines an **epimorphism** in the symplectic category as a canonical relation L for which $L \circ L^* = \text{id}_P$. From these definitions we see that a canonical relation is an isomorphism if and only if it is the graph of a symplectomorphism.

Canonical lifts of relations

The full subcategory of \mathfrak{S} consisting of cotangent bundles possesses some special properties due to the Schwartz transform $S_{M,N} : T^*N \times \overline{T^*M} \rightarrow T^*(M \times N)$, which enables us to identify canonical relations in $\text{Hom}(T^*M, T^*N)$ with immersed lagrangian submanifolds in $T^*(M \times N)$.

Example 5.29 A smooth relation between two manifolds M, N is a smooth submanifold S of the product $M \times N$. Under the Schwartz transform, the conormal bundle of S identifies with a canonical relation $L_S \in \text{Hom}(T^*N, T^*M)$ called the **cotangent lift** of S . In view of Example 3.27, this definition generalizes tangent lifts of diffeomorphisms. Note in particular

⁸Another point of view is to *define* $\text{Hom}(P, Q)$, not as a set, but rather as the object $Q \times \overline{P}$ in the symplectic category; then the composition $\text{Hom}(P, Q) \times \text{Hom}(Q, R) \rightarrow \text{Hom}(P, R)$ is the canonical relation which is the product of three diagonals.

that although a smooth map $f : M \rightarrow N$ does not in general give rise to a well-defined transformation $T^*N \rightarrow T^*M$, its graph in $M \times N$ nevertheless generates a canonical relation $L_f \in \text{Hom}(T^*N, T^*M)$.

Example 5.30 The diagonal embedding $\Delta : M \rightarrow M \times M$ induces a canonical relation $L_\Delta \in \text{Hom}(T^*(M \times M), T^*M)$ which in local coordinates assumes the form

$$L_\Delta = \{((x, \alpha), (x, x, \beta, \gamma)) : x \in M \text{ and } \alpha, \beta, \gamma \in T_x^*M \text{ satisfy } \alpha = \beta + \gamma\}.$$

In other words, L_Δ is the graph of addition in the cotangent bundle. If $L_1, L_2 \in \text{Hom}(*, T^*M)$ and $L_2 \times L_1$ is identified with an element of $\text{Hom}(*, T^*(M \times M))$ by the usual symplectomorphism $T^*M \times T^*M \rightarrow T^*(M \times M)$, then the **sum** $L_1 + L_2 \in \text{Hom}(*, T^*M)$ is defined as

$$L_1 + L_2 = L_\Delta \circ (L_2 \times L_1).$$

Note that if L_1, L_2 coincide with the images of closed 1-forms φ_1, φ_2 on M , then $L_1 + L_2$ equals the image of $\varphi_1 + \varphi_2$.

The dual of the isomorphism $T^*M \times T^*M \rightarrow T^*(M \times M)$ identifies L_Δ with an element $L'_\Delta \in \text{Hom}(T^*M, \text{Hom}(T^*M, T^*M))$. This canonical relation satisfies $L_1 + L_2 = L'_\Delta(L_1)(L_2)$, and if L is the image of a closed 1-form φ on M , then

$$L'_\Delta(L) = f_\varphi,$$

where f_φ is the fiberwise translation mapping introduced in Section 3.2.

△

Example 5.31 (The Legendre transform) As a particular example of this situation, let V be a smooth manifold, and consider the fiber product $TV \times_V T^*V$ along with its natural inclusion $\iota : TV \times_V T^*V \hookrightarrow TV \times T^*V$ and “evaluation” function $ev : TV \times_V T^*V \rightarrow \mathbb{R}$ given by

$$ev((x, v), (x, p)) = \langle p, v \rangle.$$

If $L \subset T^*(TV \times_V T^*V)$ is the lagrangian submanifold given by the image of the differential $d(ev)$, then the push-forward $L_{leg} = \iota_*L$ defines an isomorphism in $\text{Hom}(T^*(TM), T^*(T^*M))$ given in local coordinates by

$$((x, v), (\xi, \eta)) \mapsto ((x, \eta), (\xi, -v)).$$

As noted in [57], this canonical relation can be viewed as a geometric representative of the **Legendre transform** in the following sense: If $L : TM \rightarrow \mathbb{R}$ is a hyperregular Lagrangian function, in the sense that its fiber-derivative defines a diffeomorphism $TM \rightarrow T^*M$, then the composition of L_{leg} with the image of dL equals the image of dH , where $H : T^*M \rightarrow \mathbb{R}$ is the classical Legendre transform of L .

△

Morphisms associated with coisotropic submanifolds

If $C \subset P$ is a reducible coisotropic submanifold, then the **reduction relation** $R_C \in \text{Hom}(P, C/\mathcal{C}^\perp)$ is defined as the composition of the quotient map $C \rightarrow C/\mathcal{C}^\perp$ and the inclusion $C/\mathcal{C}^\perp \rightarrow P$. Somewhat more concretely, the relation R_C is the subset

$$\{([x], x) : x \in C\}$$

of $C/\mathcal{C}^\perp \times P$, where $[x]$ is the leaf of the characteristic foliation passing through $x \in C$. Evidently R_C is an epimorphism, and so $R_C \circ R_C^* = \text{id}_P$. On the other hand, R_C is not a monomorphism unless $C = P$, and we define the **projection relation** $K_C \in \text{Hom}(P, P)$ as the composition $R_C^* \circ R_C$, i.e.

$$K_C = \{(x, y) : x, y \in C, [x] = [y]\}.$$

By associativity, we have $K_C = K_C \circ K_C$, and $K_C^* = K_C$. Thus, K_C is like an orthogonal projection operator.

These relations give a simple interpretation of Theorem 5.12 in the symplectic category. If L is an immersed lagrangian submanifold of P , i.e. $L \in \text{Hom}(Z, P)$, then L is composable with both R_C and K_C , provided that it intersects C cleanly. In this case, we have

$$L_C = R_C \circ L \qquad L^C = K_C \circ L.$$

In particular, K_C fixes any lagrangian submanifold of C .

Example 5.32 Suppose that \mathcal{F} is a foliation on a manifold B whose leaf space B/\mathcal{F} is smooth. In the notation of Example 5.17, we then find that if $T^*(B/\mathcal{F})$ is canonically identified with the reduced space E/\mathcal{C}^\perp , then the reduction relation R_E is the Schwartz transform of the conormal bundle to the graph of the projection $B \rightarrow B/\mathcal{F}$.

△

5.3 Symplectic manifolds and mechanics

In general, an arbitrary symplectic manifold has no associated “configuration space,” and therefore the classical and quantum mechanical viewpoints must be adapted to this new context based on the available structure.

The classical picture

The central objects in the classical picture of mechanics on an arbitrary symplectic manifold (P, ω) are the semi-classical states, represented as before by lagrangian submanifolds of P equipped with half-densities, and the vector space of observables $C^\infty(P)$.

With respect to pointwise multiplication, $C^\infty(P)$ forms a commutative associative algebra. Additionally, the symplectic form on P induces a Lie algebra structure on $C^\infty(P)$ given by the **Poisson bracket**

$$\{f, g\} = X_g \cdot f,$$

where $X_g = \tilde{\omega}^{-1}(dg)$ denotes the hamiltonian vector field associated to g . These structures satisfy the compatibility condition

$$\{fh, g\} = f\{h, g\} + \{f, g\}h,$$

and are referred to collectively by calling $C^\infty(P)$ the **Poisson algebra** of P .

The classical system evolves along the trajectories of the vector field X_H associated to the choice of a hamiltonian $H: P \rightarrow \mathbb{R}$. If f is any observable, then Hamilton's equations assume the form

$$\dot{f} = \{H, f\}.$$

Note that in local Darboux coordinates, the Poisson bracket is given by

$$\{f, g\} = \sum_j \left(\frac{\partial f}{\partial q_j} \cdot \frac{\partial g}{\partial p_j} - \frac{\partial f}{\partial p_j} \cdot \frac{\partial g}{\partial q_j} \right).$$

Setting $f = q_j$ or $f = p_j$ in the Poisson bracket form of Hamilton's equations above yields their familiar form, as in Section 3.3.

Two functions $f, g \in C^\infty(P)$ are said to be in **involution** if $\{f, g\} = 0$, in which case the hamiltonian flows of f and g commute. An observable in involution with the hamiltonian H is called a **first integral** or **constant of the motion** of the system. A collection f_i of functions in involution on P is said to be **complete** if the vanishing of $\{f_i, g\}$ for all i implies that g is a function of the form $g(x) = h(f_1(x), \dots, f_n(x))$.

The quantum mechanical picture

Quantum mechanical observables should be the vector space $\mathcal{A}(\mathfrak{H}_P)$ of self-adjoint linear operators on some complex Hilbert space \mathfrak{H}_P associated to P . The structure of a commutative, non-associative algebra is defined on \mathcal{A} by Jordan multiplication:

$$A \circ B = \frac{1}{2}(AB + BA),$$

representing the quantum analog of pointwise multiplication in the Poisson algebra of P . Similarly, the \hbar -dependent commutator

$$[A, B]_{\hbar} = \frac{i}{\hbar}(AB - BA)$$

defines a Lie algebra structure on \mathcal{A} analogous to the Poisson bracket on $C^\infty(P)$.

Quantum mechanical states are vectors in \mathfrak{H}_P . The time-evolution of the quantum system is determined by a choice of energy operator \hat{H} , which acts on states via the Schrödinger equation:

$$\dot{\psi} = \frac{i}{\hbar}\hat{H}\psi.$$

A collection $\{A_j\}$ of quantum observables is said to be **complete** if any operator B which commutes with each A_j is a multiple of the identity. This condition is equivalent to the irreducibility of $\{A_j\}$, in other words, no nontrivial subspace of \mathfrak{H}_P is invariant with respect to each j .

Quantization

Although neither an underlying configuration manifold nor its intrinsic Hilbert space is generally available in the context of arbitrary symplectic manifolds, the basic goal of quantization remains the same as in the case of cotangent bundles: starting with a symplectic manifold P , we wish to identify a $*$ -algebra \mathcal{A} of operators which give the quantum analog of the “system” at hand. Early approaches to this problem were based on the principle that, regardless of how the algebra \mathcal{A} is identified, the correspondence between classical and quantum observables should be described by a linear map from the Poisson algebra $C^\infty(P)$ to \mathcal{A} which satisfies the following criteria known as the **Dirac axioms**

Definition 5.33 *A linear map $\rho : C^\infty(P) \rightarrow S(\mathfrak{H}_P)$ is called a **quantization** provided that it satisfies*

1. $\rho(1) = \text{identity}$.
2. $\rho(\{f, g\}) = [\rho(f), \rho(g)]_{\hbar}$
3. *For some complete set of functions f_1, \dots, f_n in involution, the operators $\rho(f_1), \dots, \rho(f_n)$ form a complete commuting set.*

As was eventually proven by Groenwald and Van Hove (see [1] for a proof), a quantization of all classical observables in this sense does not exist in general.

To make the basic quantization problem more tractable, we first enlarge the class of classical objects to be quantized, but then relax the criteria by which quantum and classical objects are to correspond. Based on an idea of Weyl and von Neumann, we first replace the classical observables by the groups of which they are the infinitesimal generators. The basic classical objects are then symplectomorphisms, which should be represented by unitary operators on quantum Hilbert spaces.

A general formulation of the quantization problem is then to define a “functor” from the symplectic category to the category of (hermitian) linear spaces. This means that to each symplectic manifold P we should try to assign a Hilbert space \mathfrak{H}_P in such a way that $\mathfrak{H}_{\bar{P}}$ is dual to \mathfrak{H}_P , and $\mathfrak{H}_{P \times Q}$ is canonically isomorphic to a (completed) tensor product $\mathfrak{H}_P \widehat{\otimes} \mathfrak{H}_Q$. However this is accomplished, each canonical relation $L \in \text{Hom}(P, Q)$ must then be assigned to a linear operator $T_L \in \text{Hom}(\mathfrak{H}_P, \mathfrak{H}_Q) \simeq \mathfrak{H}_Q \otimes \mathfrak{H}_P^*$ in a way which commutes with compositions, i.e.,

$$T_{\tilde{L} \circ L} = T_{\tilde{L}} \circ T_L$$

for each $\tilde{L} \in \text{Hom}(Q, R)$. In these abstract terms, our classical and quantum correspondences can be expressed as follows.

Object	Classical version	Quantum version
basic spaces	symplectic manifold (P, ω) $Q \times P$ $Q \times \overline{P}$ point $*$	hermitian vector space \mathfrak{H}_P $\mathfrak{H}_Q \widehat{\otimes} \mathfrak{H}_P$ $\text{Hom}(\mathfrak{H}_P, \mathfrak{H}_Q)$ \mathbb{C}
state	lagrangian submanifold L	element of \mathfrak{H}_P
space of observables	Poisson algebra $C^\infty(P)$	symmetric operators on \mathfrak{H}_P

At this stage, two relevant observations are apparent from our earlier study of WKB quantization. First, in addition to just a lagrangian submanifold $L \subset P$, it may require more data (such as a symbol) to determine an element of \mathfrak{H}_P in a consistent way. Second, the Hilbert space \mathfrak{H}_P may carry some sort of filtration (e.g. by powers of \hbar or by degree of smoothness), and the quantization may be “correct” only to within a certain degree of accuracy as measured by the filtration. By “correctness” we mean that composition of canonical relations should correspond to composition of operators. As we have already seen, it is too much to require that this condition be satisfied exactly. The best we can hope for is a functorial relation, rather than a mapping, from the classical to the quantum category.

6 Fourier Integral Operators

For our first examples of a quantization theory which gives a functorial relation between the classical (symplectic) category and the quantum category of hermitian vector spaces, we return to WKB quantization in cotangent bundles. In this context, semi-classical states are represented by pairs (L, \mathfrak{s}) consisting of an immersed lagrangian submanifold L in $T^*(M \times N)$ equipped with a *symbol* \mathfrak{s} . Our discussion begins by defining a suitable notion of composition for such states. The WKB quantization $I_{\hbar}(L, \mathfrak{s})$ of the state (L, \mathfrak{s}) is then regarded as the Schwartz kernel of the corresponding operator in $\text{Hom}(\mathfrak{H}_M, \mathfrak{H}_N)$. A particular concrete case of this classical-quantum correspondence is given by the symbol calculus of Fourier integral operators.

6.1 Compositions of semi-classical states

In Section 4.4, we defined a semi-classical state in a cotangent bundle T^*M as a pair (L, \mathfrak{s}) consisting of a quantizable lagrangian submanifold $L \subset T^*M$ and a symbol \mathfrak{s} on L . In this section, we study certain natural transformations of semi-classical states.

Reduction of semi-classical states

We will say that a reducible pair (C, L) in a symplectic manifold P is **properly reducible** if the quotient of $I = L \times_P C$ by its characteristic foliation is a smooth, Hausdorff manifold and the map $I \rightarrow L_C$ is proper.

Lemma 6.1 *If (C, L) is a properly reducible pair in a symplectic manifold P , then there exists a natural linear map*

$$|\Omega|^{1/2}L \otimes |\Omega|^{1/2}C \rightarrow |\Omega|^{1/2}L_C.$$

Proof. First note that if V is a symplectic vector space, together with a lagrangian subspace L and a coisotropic subspace C , then the exact sequence

$$0 \rightarrow L \cap C^\perp \rightarrow L \cap C \rightarrow L_C \rightarrow 0$$

gives rise to the isomorphism

$$|\Lambda|^{1/2}L_C \otimes |\Lambda|^{1/2}(L \cap C^\perp) \simeq |\Lambda|^{1/2}(L \cap C).$$

The linear maps $v \mapsto (v, -v)$ and $(x, y) \mapsto x + y$ define a second exact sequence

$$0 \rightarrow L \cap C \rightarrow L \oplus C \rightarrow L + C \rightarrow 0,$$

from which we get

$$|\Lambda|^{1/2}L_C \otimes |\Lambda|^{1/2}(L \cap C^\perp) \otimes |\Lambda|^{1/2}(L + C) \simeq |\Lambda|^{1/2}L \otimes |\Lambda|^{1/2}C.$$

Finally, the exact sequence

$$0 \rightarrow (L + C)^\perp \rightarrow V \xrightarrow{\tilde{\omega}} (L + C)^* \rightarrow 0$$

combined with the half-density on V induced by the symplectic form defines a natural isomorphism

$$|\Lambda|^{1/2}(L + C) \simeq |\Lambda|^{1/2}(L + C)^\perp.$$

Since $L \cap C^\perp = (L + C)^\perp$, we arrive at

$$|\Lambda|^{1/2}L_C \otimes |\Lambda|(L \cap C^\perp) \simeq |\Lambda|^{1/2}L \otimes |\Lambda|^{1/2}C.$$

Now consider a properly reducible pair (C, L) in a symplectic manifold P . By the preceding computations, there is a linear map

$$|\Omega|^{1/2}L \otimes |\Omega|^{1/2}C \rightarrow |\Omega|^{1/2}L_C \otimes |\Omega|(\mathcal{F}_I).$$

Since the quotient map $I \rightarrow L_C$ is proper, integration over its fibers is well-defined and gives the desired linear map

$$|\Omega|^{1/2}L \otimes |\Omega|^{1/2}C \rightarrow |\Omega|^{1/2}L_C.$$

□

Let M be a smooth manifold and consider a submanifold $N \subset M$ equipped with a foliation \mathcal{F} such that the leaf space $N_{\mathcal{F}}$ is a smooth, Hausdorff manifold. From Section 5.1 we recall that the integrability of \mathcal{F} implies that

$$C_N = \{(x, p) \in T^*M : x \in N, \mathcal{F}_x \subset \ker(p)\}$$

is a coisotropic submanifold of T^*M whose reduced space is the cotangent bundle $T^*N_{\mathcal{F}}$ of the leaf space $N_{\mathcal{F}}$. If L is an immersed lagrangian submanifold of T^*M such that (C_N, L) form a reducible pair, then we denote by I the fiber product $L \times_{T^*M} C_N$ of L and C_N and consider the following commutative diagram

$$\begin{array}{ccc} L & \longrightarrow & T^*M \\ r_L \uparrow & & \uparrow \iota \\ I & \xrightarrow{r_{C_N}} & C_N \\ \pi \downarrow & & \downarrow p_C \\ L_C & \xrightarrow{J} & T^*N \end{array}$$

Our goal is describe how a symbol \mathfrak{s} on L naturally induces a symbol \mathfrak{s}_C on L_C .

Lemma 6.2 *In the notation of the diagram above, there is a natural isomorphism*

$$r_L^* \Phi_{L, \hbar} \rightarrow \pi^* \Phi_{L_C, \hbar}.$$

Proof. From the proof of Lemma 5.16 it follows easily that the pull-back to I of the Liouville forms on T^*M and T^*N coincide, and thus, the pull-back to I of the prequantum line bundles over T^*M and T^*N are canonically isomorphic.

Similarly, let E be the symplectic vector bundle over I given by the pull-back of $T(T^*M)$. Lagrangian subbundles λ, λ' of E are then induced by the immersed lagrangian submanifold L and the vertical subbundle VM of $T(T^*M)$. By definition, the pull-back of the Maslov bundle of L to I is canonically isomorphic to the bundle $M_{\lambda, \lambda'}(E)$.

Now the tangent bundle of C_N induces a coisotropic subbundle C of E , and one must check that the pull-back of the Maslov bundle of L_C to I is canonically isomorphic to $M_{\lambda_C, \lambda'_C}(C/C^\perp)$. From Theorem 5.22 we therefore obtain a canonical isomorphism of $r_L^*M_L$ with $\pi^*M_{L_C}$. By tensoring this isomorphism with the isomorphism of prequantum bundles described in the preceding paragraph, we arrive at the desired isomorphism of phase bundles.

□

A parallel section s of $\Phi_{L, \hbar}$ pulls back to a parallel section of $r_L^*\Phi_{L_C, \hbar}$, which, under the isomorphism of Lemma 6.2, identifies with a parallel section of $\pi^*\Phi_{L_C, \hbar}$. Since parallel sections of $\pi^*\Phi_{L_C, \hbar}$ identify naturally with parallel sections of $\Phi_{L_C, \hbar}$, we obtain a map

$$(\text{parallel sections of } \Phi_{L, \hbar}) \rightarrow (\text{parallel sections of } \Phi_{L_C, \hbar}).$$

By tensoring with the map of density spaces given above, we obtain, for each half-density σ on C , a natural map of symbol spaces

$$S_L \rightarrow S_{L_C}.$$

We will denote the image of a symbol \mathfrak{s} on L under this map by $\mathfrak{s}_{C, \sigma}$.

Composition of semi-classical states

If L_1, L_2 are immersed lagrangian submanifolds of T^*M, T^*N , then their product $L_2 \times L_1$ gives a well-defined immersed lagrangian submanifold of $T^*(M \times N)$ via the Schwartz transform $S_{M, N}: T^*N \times \overline{T^*M} \rightarrow T^*(M \times N)$. By the fundamental properties of the Schwartz transform described in Proposition 3.32, it follows easily that the phase bundle $\Phi_{L_2 \times L_1, \hbar}$ is canonically isomorphic to the external tensor product $\Phi_{L_2, \hbar} \boxtimes \Phi_{L_1, \hbar}^{-1}$. From the discussion in Appendix A we also have a linear isomorphism of density bundles $|\Lambda|^{1/2}(L_2 \times L_1) \rightarrow |\Lambda|^{1/2}L_2 \boxtimes |\Lambda|^{1/2}L_1$, and thus there is a natural linear map of symbol spaces

$$\mathcal{S}_{L_2} \otimes \mathcal{S}_{L_1} \rightarrow \mathcal{S}_{L_2 \times L_1}$$

which we denote by $(\mathfrak{s}_1, \mathfrak{s}_2) \rightarrow \mathfrak{s}_2 \boxtimes \mathfrak{s}_1^*$.

Definition 6.3 *The product of semi-classical states $(L_1, \mathfrak{s}_1), (L_2, \mathfrak{s}_2)$ in T^*M, T^*N is the semi-classical state $(L_2 \times L_1, \mathfrak{s}_2 \boxtimes \mathfrak{s}_1^*)$ in $T^*(M \times N)$.*

□

Turning to compositions, we first note that if M, V are smooth manifolds, then the Schwartz transform $S_{M,V} : T^*V \times \overline{T^*M} \rightarrow T^*(M \times V)$ identifies each canonical relation in $\text{Hom}(T^*M, T^*V)$ with an immersed lagrangian submanifold in $T^*(M \times V)$. Our goal is now to describe how semi-classical states (L, \mathfrak{s}) in $T^*(M \times N)$ and (L', \mathfrak{s}') in $T^*(N \times V)$ define a semi-classical state in $T^*(M \times V)$ when L, L' are composable as canonical relations in $\text{Hom}(T^*M, T^*N)$ and $\text{Hom}(T^*N, T^*V)$, respectively.

We begin with the following commutative diagram

$$\begin{array}{ccc} T^*V \times \overline{T^*N} \times T^*N \times \overline{T^*M} & \longrightarrow & T^*(M \times N \times N \times V) \\ \downarrow & & \downarrow \\ T^*V \times \overline{T^*M} & \longrightarrow & T^*(M \times V) \end{array}$$

where the upper horizontal arrow denotes the composition $S_{M \times N} \circ (S_{N,V} \times S_{M,N})$ of Schwartz transforms, and the lower horizontal arrow equals the Schwartz transform $S_{M,V}$. The left vertical arrow represents the reduction relation defined by the coisotropic submanifold $C = T^*V \times \Delta_{T^*N} \times \overline{T^*M}$, while the right vertical arrow is the reduction relation defined by the image of C under the Schwartz transform. $L' \circ L \in \text{Hom}(T^*M, T^*V)$ is defined as the reduction of $L_2 \times L_1$ by the coisotropic submanifold $T^*V \times \Delta_{T^*N} \times T^*M$ of $T^*V \times \overline{T^*N} \times T^*N \times \overline{T^*M}$. Under the Schwartz transform,

$$T^*V \times \overline{T^*N} \times T^*N \times \overline{T^*M} \rightarrow T^*(M \times N \times N \times V),$$

the submanifold $T^*V \times \Delta_{T^*N} \times T^*M$ maps to the conormal submanifold $C_{V,N,M}$ of $T^*(V \times N \times N \times M)$ defined by $V \times \Delta_N \times M \subset V \times N \times N \times M$ and its product foliation by subsets of the form $v \times \Delta_N \times m$ for $(v, m) \in V \times M$. Thus, the Schwartz transform of $L_2 \times L_1$ reduces by $C_{V,N,M}$ to yield the image of $L_2 \circ L_1$ under the Schwartz transform $T^*V \times \overline{T^*M} \rightarrow T^*(V \times M)$.

From Lemma 6.1, it follows that for any properly reducible pair L_1, L_2 , we obtain a natural linear map of symbol spaces

$$\mathcal{S}_{L_2} \otimes \mathcal{S}_{L_1} \rightarrow \mathcal{S}_{L_2 \circ L_1}$$

given by the composition of the product map above and the reduction map $\mathcal{S}_{L_2 \times L_1} \rightarrow \mathcal{S}_{L_2 \circ L_1}$ defined by the natural half-density on $C_{V,N,M}$ induced by the symplectic forms on T^*V, T^*N , and T^*M . We denote the image of $\mathfrak{s}_2 \boxtimes \mathfrak{s}_1$ under this map by $\mathfrak{s}_2 \circ \mathfrak{s}_1$.

We will say that semi-classical states (L_1, \mathfrak{s}_1) in $T^*(M \times N)$ and (L_2, \mathfrak{s}_2) in $T^*(N \times V)$ are **composable** if the conormal submanifold $C_{V,N,M}$ and immersed lagrangian submanifold $L_2 \times L_1$ form a properly reducible pair.

Definition 6.4 *If a semi-classical state (L_1, \mathfrak{s}_1) in $T^*(M \times N)$ is composable with a semi-classical state (L_2, \mathfrak{s}_2) in $T^*(N \times V)$, then their composition is defined as the semi-classical state $(L_2 \circ L_1, \mathfrak{s}_2 \circ \mathfrak{s}_1)$ in $T^*(M \times V)$.*

□

Example 6.5 Consider a semi-classical state $(L_F, \tilde{\mathfrak{s}})$ in $T^*(M \times N)$, where L_F is the Schwartz transform of the graph of a symplectomorphism $F: T^*M \rightarrow T^*N$. If L is an immersed lagrangian submanifold of T^*M , then Definition 6.4 provides that for each symbol $\mathfrak{s} \in \mathcal{S}_L$, the semi-classical state (L, \mathfrak{s}) in T^*M is transformed by F into the semi-classical state $(L_F \circ L, \tilde{\mathfrak{s}} \circ \mathfrak{s})$ in T^*N .

To make this correspondence more explicit, let us choose a particular lagrangian immersion $\iota: L \rightarrow T^*M$ representing L (see the discussion in Section 4.4). We then note that the natural half-density $|\omega_M^n|^{1/2}$ on $L_F \approx T^*M$ enables us to identify the symbol space \mathcal{S}_{L_F} with the product

$$\Gamma_{L_F} \otimes_{\mathbb{C}} C^\infty(L_F, \mathbb{C}).$$

Given a parallel section \tilde{s} and smooth complex-valued function h on L_F , we find by a computation that the isomorphism $\mathcal{S}_L \rightarrow \mathcal{S}_{L_F \circ L}$ determined by the symbol $\tilde{\mathfrak{s}} = \tilde{s} \otimes h$ is given by $s \otimes a \mapsto s' \otimes (a \cdot \iota^* h)$, where a is any half-density on L and s' is the unique parallel section of $\Phi_{L_F \circ L, \tilde{h}}$ such that $\iota^* \tilde{s} \mapsto s' \otimes s^{-1}$ under the canonical isomorphism $\iota^* \Phi_{L_F, \tilde{h}} \simeq \Phi_{L_F \circ L, \tilde{h}} \otimes \Phi_{L, \tilde{h}}^{-1}$.

△

Example 6.6 As a special case of Example 6.5, note that if the symplectomorphism $F: T^*M \rightarrow T^*N$ is the cotangent lift of a diffeomorphism $M \rightarrow N$, then for each $\hbar > 0$, the phase bundle of L_F admits a canonical parallel section, and so the symbol space \mathcal{S}_{L_F} identifies naturally with $\mathbb{R}_+ \times C^\infty(L_F, \mathbb{C})$. If $h \in C^\infty(L_F, \mathbb{C})$, then with this identification, the composition of (L_F, h) with a semi-classical state (L, \mathfrak{s}) equals the semi-classical state $(L_F \circ L, (h \circ \iota) \cdot \mathfrak{s})$.

Similarly, if β is a closed 1-form on M and the symplectomorphism $F: T^*M \rightarrow T^*M$ equals fiberwise translation by β , then $\hbar \in \mathbb{R}_+$ is admissible for L_F if and only if $[\beta]$ is \hbar -integral. For such \hbar , each parallel section of the phase bundle $\Phi_{L_F, \hbar}$ identifies with an oscillatory function of the form $ce^{iS/\hbar}$ for $c \in \mathbb{R}$ and some $S: M \rightarrow \mathbb{T}_\hbar$ satisfying $S^* d\sigma = \beta$. A computation then shows that the composition of a semi-classical state of the form $(L_F, e^{iS/\hbar} \otimes h)$ with (L, ι, \mathfrak{s}) yields the semi-classical state $(L_F \circ L, e^{i(S \circ \pi_L)/\hbar} (h \circ \iota) \cdot \mathfrak{s})$.

△

6.2 WKB quantization and compositions

To define the composition of semi-classical states $(L, \mathfrak{s}) \in \text{Hom}(T^*M, T^*N)$ and $(\tilde{L}, \tilde{\mathfrak{s}}) \in \text{Hom}(T^*N, T^*V)$ as a semi-classical state $(\tilde{L} \circ L, \tilde{\mathfrak{s}} \circ \mathfrak{s}) \in \text{Hom}(T^*M, T^*V)$, we used the Schwartz transform to identify the immersed lagrangian submanifold $\tilde{L} \circ L \in T^*V \times \overline{T^*M}$ with an immersed lagrangian submanifold in $T^*(M \times V)$, since it is the cotangent bundle structure of the latter space which gives meaning to “symbols” (in the sense of Chapter 4) on $\tilde{L} \circ L$. On the quantum level, an analogous correspondence is furnished by the Schwartz kernel theorem.

Let M be a smooth manifold and let $|\Omega|_0^{1/2} M$ be equipped with the topology of C^∞ convergence on compact sets. A **distributional half-density** on M is then a continuous, \mathbb{C} -linear functional on $|\Omega|_0^{1/2} M$. We denote the space of distributional half-densities on M

by $|\Omega|_{-\infty}^{1/2}M$, and we equip this space with the weak* topology. If N is another smooth manifold, a **kernel** is any element of $|\Omega|_{-\infty}^{1/2}(M \times N)$. A kernel K defines a linear map $\mathcal{K}: |\Omega|_0^{1/2}M \rightarrow |\Omega|_{-\infty}^{1/2}N$ by duality via the equation

$$\langle \mathcal{K}(u), v \rangle \stackrel{def}{=} \langle K, u \otimes v \rangle \quad (*)$$

Schwartz kernel theorem . *Every $K \in |\Omega|_{-\infty}^{1/2}(M \times N)$ defines a linear map $\mathcal{K}: |\Omega|_0^{1/2}M \rightarrow |\Omega|_{-\infty}^{1/2}N$ by (*) above, which is continuous in the sense that $\mathcal{K}(\phi_j) \rightarrow 0$ in $|\Omega|_{-\infty}^{1/2}N$ if $\phi_j \rightarrow 0$ in $|\Omega|_0^{1/2}M$. Conversely, to every such linear map, there is precisely one distribution K such that (*) is valid.*

□

Of course, each $w \in |\Omega|_{-\infty}^{1/2}N$ defines an element $\tilde{w} \in |\Omega|_{-\infty}^{1/2}N$ by the equation

$$\langle \tilde{w}, v \rangle = \int_N v \otimes w.$$

In all of the examples we will consider, the image of the map \mathcal{K} will lie within the subspace of $|\Omega|_{-\infty}^{1/2}N$ represented by elements of $|\Omega|_0^{1/2}N$ in this way, and so we can consider \mathcal{K} as a continuous linear map $|\Omega|_0^{1/2}M \rightarrow |\Omega|_0^{1/2}N$, thereby giving a (densely defined) operator $\mathfrak{H}_M \rightarrow \mathfrak{H}_N$ on the intrinsic Hilbert spaces of M, N . In practical terms, the value of the half-density Ku at each $y \in N$ can be computed in these cases by, roughly speaking, integrating the product $K \otimes u$ over the submanifold $\{y\} \times M$ of $N \times M$. Thus, in the same way that the Schwartz transform provides a natural correspondence

$$\text{Hom}(T^*M, T^*N) \leftrightarrow T^*(M \times N),$$

the Schwartz kernel theorem gives the identification

$$|\Omega|_{-\infty}^{1/2}(M \times N) \leftrightarrow \text{Hom}(\mathfrak{H}_M, \mathfrak{H}_N)$$

\hbar -differential operators

With respect to linear coordinates $\{x_j\}$ on \mathbb{R}^n , we define \hbar -dependent operators

$$D_j \stackrel{def}{=} -i\hbar \frac{\partial}{\partial x_j}.$$

An \hbar -differential operator of asymptotic order⁹ $k \in \mathbb{Z}$ is then an asymptotic series of the form

$$P_{\hbar} = \sum_{m=0}^{\infty} P_m \hbar^{m+k},$$

⁹Note that this order generally differs from the order of a differential operator.

where each P_m is a polynomial in the operators D_j . By formally substituting ξ_j for D_j in each term of P_{\hbar} , we obtain an \hbar -dependent function $\sigma_{P_{\hbar}}$ on $T^*\mathbb{R}^n$, known as the **symbol** of P_{\hbar} , which is related to P_{\hbar} by the asymptotic Fourier inversion formula: for any compactly supported oscillatory test function $e^{-i\psi/\hbar}u$ (we permit $\psi = 0$), we have

$$(P_{\hbar}e^{-i\psi/\hbar}u)(x) = (2\pi\hbar)^{-n} \iint e^{i\langle(x-y,\xi)-\psi\rangle/\hbar} \sigma_{P_{\hbar}}(x, \xi) u(y) dy d\xi.$$

The **principal symbol** p_{\hbar} of the operator P_{\hbar} is defined as the symbol of the first nonvanishing term in its series expansion, i.e.,

$$p_{\hbar}(x, \xi) = \sigma_{P_{m_0}}(x, \xi) \hbar^{m_0+k}$$

if $P_m = 0$ for all $m < m_0$ and $P_{m_0} \neq 0$. By replacing $\sigma_{P_{\hbar}}$ by p_{\hbar} in the formula above, we obtain

$$(P_{\hbar}e^{-i\psi/\hbar}u)(x) = (2\pi\hbar)^{-n} \iint e^{i\langle(x-y,\xi)-\psi\rangle/\hbar} p_{\hbar}(x, y, \xi) u(y) dy d\xi + O(\hbar^{m_0+k+1})$$

where p_{\hbar} is of course independent of y . (Later we will drop this assumption to obtain a more symmetric calculus). For fixed x , the function $(y, \xi) \mapsto \langle x - y, \xi \rangle - \psi$ has a nondegenerate critical point when $y = x$ and $\xi = d\psi(x)$. An application of the principle of stationary phase therefore gives

$$(P_{\hbar}e^{-i\psi/\hbar}u)(x) = e^{-i\psi(x)/\hbar}u(x) \cdot p_{\hbar}(x, x, d\psi(x)) \hbar^{m_0+k} + O(\hbar^{m_0+k+1}) \quad (**)$$

To interpret this expression in the context of WKB quantization, we first note that the phase function $\phi(x, y, \xi) = \langle x - y, \xi \rangle$ on $B = \mathbb{R}^n \times \mathbb{R}^n \times (\mathbb{R}^n)^*$, together with $V = \mathbb{R}^n \times \mathbb{R}^n$ and the cartesian projection $p_V : B \rightarrow V$ define a Morse family (B, V, p_V, ϕ) which generates the conormal bundle to the diagonal $\Delta \subset \mathbb{R}^n \times \mathbb{R}^n$. The principal symbol p_{\hbar} , written as a function of the variables (x, y, ξ) defines an amplitude $\mathfrak{a} = p_{\hbar} \cdot |dx dy|^{1/2} |d\xi|$ on B , whose restriction to the fiber-critical set $\Sigma_{\phi} = \{(x, x, \xi) : \langle x, \xi \rangle = 0\}$ induces a well-defined symbol $\mathfrak{s}_P = p_{\hbar}(x, x, \xi)$ on L_{Δ} .

Now, $e^{-i\psi/\hbar}u = I_{\hbar}(L, \mathfrak{s})$, where L is the projectable lagrangian submanifold of $T^*\mathbb{R}^n$ defined by $\text{im}(d\psi)$, and \mathfrak{s} is obtained from the pull-back of u to L . By Example 6.6, we have $(L, \mathfrak{s}) = (L_{\Delta} \circ L, \mathfrak{s}_P \circ \mathfrak{s})$, that is, composition with $(L_{\Delta}, \mathfrak{s}_P)$ multiplies the symbol \mathfrak{s} by the values of the principal symbol p_{\hbar} on $L \simeq \Sigma_{\phi}$. Combined with the preceding equation, this gives

$$(P_{\hbar}e^{-i\psi/\hbar}u)(x) = I_{\hbar}(L_{\Delta} \circ L, \mathfrak{s}_P \circ \mathfrak{s}) \cdot \hbar^{m_0+k} + O(\hbar^{m_0+k+1}).$$

The Schwartz kernel for the operator P_{\hbar} is given by the distribution family $I_{\hbar}(L_{\Delta}, \mathfrak{s}_P)$.

An \hbar -differential operator on a manifold M is an operator P_{\hbar} on $|\Omega|_0^{1/2}M$ which coincides in local coordinates with a series of the form P_{\hbar} as above. While the symbol of P_{\hbar} depends on the choice of local coordinates, its principal symbol p_{\hbar} is a well-defined function on the cotangent bundle T^*M . Using a global generating function for the conormal bundle of $\Delta \subset M \times M$ (see Example 4.29) and arguing as above, we obtain the following geometric version of $(**)$ above:

Theorem 6.7 *If (L, \mathfrak{s}) is an exact, projectable semi-classical state in T^*M and P_{\hbar} is an \hbar -differential operator of order k on M , then:*

$$P_{\hbar}(I_{\hbar}(L, \mathfrak{s})) = I_{\hbar}(L_{\Delta} \circ L, \mathfrak{s}_P \circ \mathfrak{s}) + O(\hbar^{m_0+k+1}),$$

where \mathfrak{s}_P is the symbol on L_{Δ} induced by the principal symbol p_{\hbar} of P_{\hbar} .

Roughly speaking, this theorem asserts that polynomial functions on the identity relation L_{Δ} quantize as the Schwartz kernel of differential operators on M . As usual, this correspondence is only asymptotic; many differential operators share the same principal symbol, and their actions on a given function coincide only up to terms of higher order in \hbar .

As a particular case of Theorem 6.7, the vanishing of the principal symbol p_{\hbar} on L implies that

$$P_{\hbar}(I_{\hbar}(L, \mathfrak{s})) = O(\hbar^{m_0+k+1}),$$

i.e., that $I_{\hbar}(L, \mathfrak{s})$ is a first-order approximate solution to the equation $P_{\hbar}u = 0$. This remark suggests the quantum analog of certain coisotropic submanifolds of T^*M . The zero set of the principal symbol p_{\hbar} is called the **characteristic variety** of the operator P_{\hbar} . If 0 is a regular value of p_{\hbar} , then $C_{P_{\hbar}} = p_{\hbar}^{-1}(0)$ is a coisotropic submanifold of T^*M , and semi-classical states contained in $C_{P_{\hbar}}$ represent solutions of the asymptotic differential equation $P_{\hbar}u = 0$. In this sense, $C_{P_{\hbar}}$, or, more properly, the reduced manifold of $C_{P_{\hbar}}$, corresponds to the kernel of P_{\hbar} in \mathfrak{H}_M , and the projection relation $K_{C_{P_{\hbar}}}$ quantizes as the orthogonal projection onto this subspace.

A familiar illustration of these concepts is provided by the WKB approximation.

Example 6.8 Recall that the Schrödinger operator associated to a given potential V on a riemannian manifold M is given by

$$\hat{H} = -\frac{\hbar^2}{2m}\Delta + m_V.$$

For $E > 0$, the time-independent Schrödinger equation is then

$$P_{\hbar}\psi = 0,$$

where $P_{\hbar} = \hat{H} - E$ is the zeroth-order asymptotic differential operator on M with principal symbol

$$p_{\hbar}(x, \theta) = -\frac{|\theta|^2}{2m} + (V(x) - E).$$

The characteristic variety of p_{\hbar} is simply the level set $H^{-1}(E)$ of the classical hamiltonian of the system, and as in previous sections we see that first-order approximate solutions to the time-independent Schrödinger equation arise from semi-classical states represented by quantizable lagrangian submanifolds in $C_{P_{\hbar}}$.

△

The difficulty of quantizing more general symbols on the identity relation L_Δ lies in the convergence of the integral

$$(P_\hbar u)(x) = (2\pi\hbar)^n \int \int e^{i\langle x-y, \theta \rangle / \hbar} a(x, y, \xi) u(y) d\xi dy.$$

For differential operators, integration over the phase variables θ was well-defined due to the fact that the symbol of a differential operator has a polynomial growth rate with respect to these variables. Weakening this condition while still guaranteeing that the integral converges leads to the definition of **pseudodifferential operators**.

In \mathbb{R}^n , an \hbar -pseudodifferential operator of order μ is given by

$$(Au)(x) = (2\pi\hbar)^n \int \int e^{i\langle x-y, \theta \rangle / \hbar} a_\hbar(x, \theta) u(y) d\theta dy,$$

where the symbol $a(x, \theta)$ is an asymptotic series in \hbar of the form

$$a_\hbar(x, \theta) \sim \sum a_j(x, \theta) \hbar^{\mu+j}.$$

Each coefficient a_j is a smooth function on $\mathbb{R}^n \times \mathbb{R}^n \setminus \{0\}$ and is positively homogeneous of degree $\mu - j$, i.e.,

$$a_j(x, c\theta) = c^{\mu-j} a_j(x, \theta)$$

for $c > 0$. As in the special case of differential operators, the invariance of pseudodifferential operators under coordinate changes enables one to extend this theory to any smooth manifold M . In this case, the principal symbol of a pseudodifferential operator is a well-defined function on T^*M , and there is a corresponding version of Theorem 6.7 (see for example [51]).

Fourier integral operators

In its simplest form, the theory of Fourier integral operators provides a means for quantizing half-densities on more general lagrangian submanifolds L of T^*M by replacing the function $(x, y, \xi) \mapsto \langle x - y, \xi \rangle$ in the definition of pseudodifferential operators by a phase function which generates L :

$$(Au)(x) = (2\pi\hbar)^n \int \int e^{i\phi(x, y, \xi) / \hbar} a(x, y, \xi) u(y) d\xi dy.$$

As before, the class of quantizable half-densities (and canonical relations) is constrained by the necessary conditions for this integral to converge. In this section, we will only be concerned with a few specialized cases; a detailed description of this theory can be found in [21, 28, 31, 43, 56].

Perhaps the simplest generalization of the picture of (pseudo)differential operators given above is provided by quantizing semi-classical states whose underlying lagrangian submanifolds are conormal bundles. Let M be a smooth manifold with a smooth submanifold N . As constructed in Example 4.29, the conormal bundle $L_N \subset T^*N$ is generated by a single

Morse family (B, V, p_V, ϕ) with the properties that V is a tubular neighborhood of N in M and B is a vector bundle over V with fiber dimension equal to the codimension of N in M . As in the case of (pseudo)differential operators, an amplitude \mathfrak{a} on B having an asymptotic expansion in terms which are positively homogeneous with respect to the natural \mathbb{R}_+ action on the fibers of $p_V : B \rightarrow V$ gives rise to a well-defined distributional half-density family $I_{\hbar}(L, \mathfrak{s})$ on M .

It is easy to see that the fiber-critical set Σ_ϕ is invariant under the \mathbb{R}_+ action on B , and that the identification $L_N \simeq \Sigma_\phi$ defined by λ_ϕ is equivariant with respect to the natural \mathbb{R}_+ action on L_N . Consequently, positively homogeneous symbols on L_N induce positively homogeneous amplitudes on B of the same order, and we can proceed to define $I_{\hbar}(L, \mathfrak{s})$ as before by requiring that \mathfrak{s} be homogeneous.

When $M = X \times Y$ is a product manifold, the distributions $I_{\hbar}(L, \mathfrak{s})$ represent Schwartz kernels for continuous linear operators $|\Omega|_0^{1/2}X \rightarrow |\Omega|_\infty^{1/2}Y$. Under certain additional restrictions on N which are always fulfilled, for example, when N is the graph of a diffeomorphism $X \rightarrow Y$, these operators map $|\Omega|_0^{1/2}X$ to $|\Omega|_0^{1/2}Y$ and extend continuously to $|\Omega|_{-\infty}^{1/2}X$. Moreover, they satisfy the composition law

$$I_{\hbar}(\tilde{L}, \tilde{\mathfrak{s}}) \circ I_{\hbar}(L, \mathfrak{s}) = I_{\hbar}(\tilde{L} \circ L, \tilde{\mathfrak{s}} \circ \mathfrak{s}).$$

Example 6.9 If X, Y are smooth manifolds and $f : X \rightarrow Y$ is a smooth diffeomorphism, then the procedure described above can be applied to the graph Γ_f of f , viewed as a smooth submanifold of $X \times Y$. The family $I_{\hbar}(L_f, \mathfrak{s})$ of distributions on $X \times Y$ defined above corresponds via the Schwartz kernel theorem to a family of continuous linear maps $|\Omega|_0^{1/2}X \rightarrow |\Omega|_{-\infty}^{1/2}Y$

A diffeomorphism $f : X \rightarrow Y$ induces a unitary operator on intrinsic Hilbert spaces given by pull-back:

$$(f^*)^{-1} : \mathfrak{H}_X \rightarrow \mathfrak{H}_Y.$$

At the same time, f gives rise to a symplectomorphism of cotangent bundles:

$$(T^*f)^{-1} : T^*X \rightarrow T^*Y.$$

Of course, the more interesting Fourier integral operators from $|\Omega|^{1/2}X$ to $|\Omega|^{1/2}Y$ come from quantizing canonical relations from T^*X to T^*Y which do not arise from diffeomorphisms from X to Y . For instance, quantizing the hamiltonian flow $\{\varphi_t\}$ of a hamiltonian H on T^*X gives the solution operators $\exp(-it\hat{H}/\hbar)$ of the Schrödinger equation.

△

7 Geometric Quantization

7.1 Prequantization

In the preceding chapter, we observed that the characteristic variety $C = C_{P_{\hbar}}$ of a differential operator P_{\hbar} on a riemannian manifold M represents the classical analog of the kernel of P_{\hbar} in \mathfrak{H}_M , in the sense that semi-classical states contained in C quantize to first-order approximate solutions of the equation $P_{\hbar} = 0$. When C is reducible, the lagrangian submanifolds of M contained in C correspond to lagrangian submanifolds of the reduced manifold C/\mathcal{C}^{\perp} . This suggests that if C is reducible, then a quantum Hilbert space \mathfrak{H}_C somehow associated to the reduced manifold C/\mathcal{C}^{\perp} should map isometrically onto the kernel of P_{\hbar} via an appropriate quantization of the adjoint reduction relation $R_C^* \in \text{Hom}(C/\mathcal{C}^{\perp}, T^*M)$. Before turning to a systematic means for defining \mathfrak{H}_C , let us consider what conditions the reduced manifold C/\mathcal{C}^{\perp} must satisfy in order that quantizations of the reduction and projection relations be well-defined for arbitrary quantizable semi-classical states. Specifically, we ask what assumptions on C guarantee that the classical projection operation K_C preserves the class of quantizable lagrangian submanifolds in T^*M .

For the time being, we will ignore the Maslov correction and assume that C/\mathcal{C}^{\perp} is simply-connected. A simple argument then shows that K_C preserves the class of prequantizable lagrangian submanifolds provided that the Liouville class of each \mathcal{C}^{\perp} leaf is \hbar -integral.

To interpret this condition in terms of C/\mathcal{C}^{\perp} , first note that if C/\mathcal{C}^{\perp} is simply-connected, any class $[a] \in H_2(C/\mathcal{C}^{\perp}; \mathbb{Z})$ is represented by a continuous map

$$(D, \partial D) \xrightarrow{f} (C/\mathcal{C}^{\perp}, [I])$$

for a fixed leaf I of \mathcal{C}^{\perp} . By the homotopy lifting property, f lifts to a continuous map

$$(D, \partial D) \xrightarrow{\tilde{f}} (C, I).$$

Applying Stokes' theorem, we then have

$$\int_{[a]} \omega_C = \int_{\partial D} \tilde{f}^* \alpha_M \in \mathbb{Z}_{\hbar}$$

by the assumption that the Liouville class of I is \hbar -integral. Thus, we conclude that K_C preserves prequantizable lagrangian submanifolds provided that the reduced symplectic form ω_C is itself \hbar -integral. In this case, the reduced manifold C/\mathcal{C}^{\perp} is said to be **prequantizable**.

Example 7.1 From Example 5.13 we recall that the standard metric on the unit n -sphere S^n induces a kinetic energy function k_n whose constant energy surfaces C_E are reducible coisotropic submanifolds of T^*S^n for $E > 0$. Each leaf of the characteristic foliation of C_E is a circle S which projects diffeomorphically to a great circle in S^n . If γ parametrizes a such a geodesic, then its lift to a parametrization of S satisfies $\dot{\gamma} = (2E)^{-1/2} X_{k_n}$. Since $\alpha_M(X_{k_n}) = 2E$, the Liouville class of S is determined by the number

$$\int_S \alpha_M = 2\pi \cdot (2E)^{1/2}.$$

Thus C_E is prequantizable if and only if $E = (n\hbar)^2/2$ (compare Example 4.2).

△

By Theorem D.2, this condition is equivalent to the existence of a principal \mathbb{T}_\hbar bundle Q over C/C^\perp . In fact, the bundle Q can be constructed explicitly from the prequantum \mathbb{T}_\hbar bundle Q_M over T^*M by noting that for any leaf I of C^\perp , the mod- \mathbb{Z}_\hbar reduction of λ_I represents the holonomy of $Q_M|_I$; if λ_I is \hbar -integral, there exists a parallel section of Q_M over I . If C/C^\perp is prequantizable, parallel sections over the leaves of C^\perp define a foliation \mathcal{Q} of $Q_M|_L$ whose leaf space is a principal \mathbb{T}_\hbar bundle Q over C/C^\perp . The connection on Q induced by the connection on Q_M has curvature equal to the reduced symplectic form ω_C .

A linear mapping $\rho : C^\infty(P) \rightarrow S(\mathfrak{H}_P)$ satisfying the first two Dirac axioms will be called a **prequantization** of the classical system represented by P . The basic example of prequantization is that of cotangent bundles due to Segal [52], Koopman, and Van Hove. In this case, the prequantum Hilbert space is taken to be the completion of the space of smooth, complex-valued functions on the cotangent bundle $P = T^*M$ itself, with respect to the inner-product

$$\langle f, g \rangle = \int_P f \bar{g} \omega_M^n.$$

The mapping $C^\infty(P) \rightarrow S(\mathfrak{H}_P)$ is then given explicitly by the formula

$$\rho(f) = -i\hbar X_f + m_{L(f)},$$

where as usual $m_{L(f)}$ denotes multiplication by the lagrangian $L(f) = f - \alpha_M(X_f)$. Clearly the map ρ is linear and satisfies the first Dirac axiom; verification of the second is also straightforward and will be carried out in somewhat more generality below.

The Segal prequantization is an important first step towards attempts at (pre)-quantizing more general symplectic manifolds. Roughly speaking, the idea is the following: Although the symplectic form of an arbitrary symplectic manifold (P, ω) is not exact, we can choose a covering of P by open sets U_j on which the restriction of the symplectic form satisfies $\omega = -d\alpha_j$ for appropriately chosen 1-forms α_j on U_j . In direct analogy with the Segal prequantization, we can then associate to a function $f \in C^\infty(P)$ the operator

$$\rho(f)_j = -i\hbar X_f + m_{L_j(f)},$$

on $C^\infty(U_j)$. In order to associate a “global” operator to f , we hope to piece these local operators together. One way of doing this would be to impose the condition that each $\rho(f)_j$ and $\rho(f)_k$ coincide as operators on the function space $C^\infty(U_j \cap U_k)$. This essentially requires that the 1-forms α_j agree on overlaps, and we arrive at nothing new.

A true generalization of the Segal construction due to Kostant and Souriau is achieved by first reinterpreting $\rho(f)$ as an operator on sections of a line bundle over $P = T^*M$. More precisely, let E be the complex line bundle over P associated to the trivial principal \mathbb{T}_\hbar bundle $Q = P \times \mathbb{T}_\hbar$ via the representation $x \mapsto e^{-ix/\hbar}$ of \mathbb{T}_\hbar in $U(1)$. The space of sections of E identifies canonically with $C^\infty(P)$ by means of the constant section $s = 1$, and we have

$$\rho(f) = -i\hbar \nabla_{X_f} + m_f,$$

where ∇ denotes the connection on E induced by the connection 1-form $\varphi = -\alpha_M + d\sigma$ on Q . Returning to the general case, we find that the operators $\rho(f)_j$ enjoy a similar interpretation provided that $-\alpha_j + d\sigma$ defines a local representative of a connection 1-form on a (possibly nontrivial) principal \mathbb{T}_\hbar bundle Q over P . The curvature of such a bundle necessarily coincides with the symplectic form on P ; according to the discussion in Appendix C, this condition can be satisfied if and only if ω is \hbar -integral.

Definition 7.2 *A prequantization of a symplectic manifold (P, ω) is a principal \mathbb{T}_\hbar bundle Q over P equipped with a connection 1-form φ having curvature ω .*

The upshot of the Kostant-Souriau construction is that prequantizable symplectic manifolds have prequantizable Poisson algebras. While a direct proof of this fact follows the outline of the preceding paragraph, we will study prequantizable manifolds in somewhat more detail below. The prequantum Hilbert space in this case will be the completion of the vector space of smooth sections of a hermitian line bundle associated to Q . Although several more modifications of this choice must be made in order to arrive at a reasonable substitute for the intrinsic Hilbert space of the base of a cotangent bundle, this is an important first step in our general quantization program.

For the remainder of this section, we will focus on geometric properties of prequantum circle bundles and prove that their existence coincides with the prequantizability of $C^\infty(P)$.

Automorphisms of (Q, φ)

Let P be a prequantizable symplectic manifold with prequantum \mathbb{T}_\hbar bundle Q and connection φ . To a function $f \in C^\infty(P)$, we associate an operator on $C^\infty(Q)$:

$$\xi_f = \overline{X}_f - fX,$$

where \overline{X}_f denotes the horizontal lift of the hamiltonian vector field of f , and X is the fundamental vector field on Q defined by the equations

$$X \lrcorner \varphi = 1 \quad X \lrcorner d\varphi = 0.$$

A direct computation shows that the connection form φ is invariant under the flow of $\xi_f^{(k)}$. Conversely, if $\mathcal{L}_\xi \varphi = 0$ for some vector field ξ on Q , then we can decompose ξ into its horizontal and vertical parts:

$$\xi = \overline{\xi} - gX$$

for some real-valued function g on Q satisfying $dg = \overline{\xi} \lrcorner d\varphi$. From the definition of X it follows that $X \cdot g = 0$ and $[\xi, X] = 0$, and consequently, $[\overline{\xi}, X] = 0$. Thus, $\overline{\xi}$ is the horizontal lift of X_g , and so $\xi = \xi_g$. Moreover, the requirement that the curvature of φ equal the symplectic form ω implies

$$[\xi_f, \xi_g] = \overline{[X_f, X_g]} + \omega(X_f, X_g)X - 2\{f, g\}X = \xi_{\{f, g\}}.$$

The association $f \mapsto \xi_f$ therefore defines a Lie algebra isomorphism between the Poisson algebra $C^\infty(P)$ and the space $\chi(Q, \varphi)$ of φ -preserving vector fields on Q with the standard Lie bracket. This produces the exact sequence of Lie algebras

$$0 \rightarrow \mathbb{R} \rightarrow \chi(Q, \varphi) \rightarrow \chi(P, \omega) \rightarrow H^1(P; \mathbb{R}) \rightarrow 0,$$

where $H^1(P; \mathbb{R})$ is assigned the trivial bracket, and the image of $\chi(Q, \varphi)$ in $\chi(P, \omega)$ consists precisely of the hamiltonian vector fields on P . This sequence can be integrated to give an exact sequence of automorphism groups as follows. Let (Q, φ) be a prequantization of (P, ω) , and let $\text{Aut}(Q, \varphi), \text{Aut}(P, \omega)$ denote those groups of diffeomorphisms which preserve φ, ω respectively. By the definition of X it follows that every $F \in \text{Aut}(Q, \varphi)$ preserves X and is therefore \mathbb{T}_\hbar -equivariant. In particular, this means that F is the lift of a diffeomorphism f of P ; from the fact that $\pi^*\omega = d\varphi$ it follows furthermore that $f^*\omega = \omega$, that is, $f \in \text{Aut}(P, \omega)$. The association $F \mapsto f$ defines a group homomorphism

$$\text{Aut}(Q, \varphi) \rightarrow \text{Aut}(P, \omega).$$

To determine its kernel, we simply note that the identity map on P is covered by precisely those automorphisms of Q given by the action of elements of \mathbb{T}_\hbar on its connected components. This implies that the kernel is isomorphic to $H^0(P, \mathbb{T}_\hbar)$; if P (and hence Q) is connected, this is just the circle \mathbb{T}_\hbar and we have the exact sequence

$$0 \rightarrow \mathbb{T}_\hbar \rightarrow \text{Aut}(Q, \varphi) \rightarrow \text{Aut}(P, \omega),$$

i.e. $\text{Aut}(Q, \varphi)$ is a central extension of $\text{Aut}(P, \omega)$ by \mathbb{T}_\hbar . We interpret this observation to mean that an automorphism of Q is determined “up to phase” by an automorphism of P .

If we equip the space $C^\infty(Q)$ of *complex-valued* functions on Q with the inner-product

$$\langle u, v \rangle = \int_Q u \bar{v} \mu,$$

where μ denotes the volume form $\varphi \wedge (d\varphi)^n$ on Q , then each $F \in \text{Aut}(Q, \varphi)$ preserves μ , and therefore defines a unitary operator U_F on $C^\infty(Q)$ by composition:

$$U_F(u) = u \circ F.$$

Clearly the correspondence $F \mapsto U_F$ defines a unitary representation of $\text{Aut}(Q, \varphi)$ on $L^2(Q)$; from the exact sequence above, we therefore obtain a *projective* unitary representation of the image of $\text{Aut}(Q, \varphi)$ in $\text{Aut}(P, \omega)$.

Example 7.3 Translations of \mathbb{R}^{2n} are generated by the hamiltonian vector fields associated to linear functionals on \mathbb{R}^{2n} . A basis for this space is given by the vector fields

$$X_{q_i} = -\frac{\partial}{\partial p_i} \quad X_{p_i} = \frac{\partial}{\partial q_i},$$

which assume the form

$$\xi_{q_i} = -\frac{\partial}{\partial p_i} - q_i X \quad \xi_{p_i} = \frac{\partial}{\partial q_i}$$

when lifted to $Q = \mathbb{R}^{2n} \times \mathbb{T}_\hbar$ via the Segal prescription $\xi_f = \overline{X}_f - L(f)X$, where again $L(f) = f - \alpha_M(X_f)$. Our earlier results show that $[\xi_{q_i}, \xi_{p_j}] = \xi_{\{q_i, p_j\}} = \delta_{ij}X$, and so the vector fields ξ_{q_i}, ξ_{p_i} generate a Lie subalgebra $\mathfrak{h}_n \subset \chi(Q, \varphi)$ isomorphic to $\mathbb{R}^{2n} \times \mathbb{R}$ with bracket given by

$$[(v, a), (w, b)] = (0, \omega(v, w)).$$

By exponentiating, we find that \mathfrak{h}_n corresponds to a subgroup $\mathfrak{H}_n \subset \text{Aut}(Q, \varphi)$ which comprises a central extension of the translation group:

$$0 \rightarrow \mathbb{T}_\hbar \rightarrow \mathcal{H}_n \rightarrow \mathbb{R}^{2n} \rightarrow 0.$$

The group \mathcal{H}_n is known as the **Heisenberg group** of \mathbb{R}^{2n} with its usual symplectic structure. In explicit terms, \mathcal{H}_n is diffeomorphic to $\mathbb{R}^{2n} \times \mathbb{T}_\hbar$, with group multiplication given by

$$(Q, P, \sigma) \cdot (Q', P', \sigma') = (Q + Q', P + P', \sigma') = (Q + Q', P + P', \sigma + \sigma' + \sum_j P_j \cdot Q'_j).$$

△

Kostant-Souriau prequantization

To prequantize the Poisson algebra $C^\infty(P)$, we first recall that complex line bundles E_k associated to Q arise via representations of \mathbb{T}_\hbar in $U(1)$ of the form $x \mapsto e^{ikx/\hbar}$. Smooth sections of E_k identify with functions on Q satisfying the equivariance condition

$$f(p \cdot a) = e^{-ika/\hbar} f(p)$$

for $a \in \mathbb{T}_\hbar$; in other words, the space of sections of E_k is isomorphic to the $-ik/\hbar$ -eigenspace \mathcal{E}_k of the fundamental vector field X . Under this correspondence, covariant differentiation by a vector field η on P is given simply by the Lie derivative with respect to the horizontal lift $\bar{\eta}$ of η to Q .

To each $f \in C^\infty(P)$ and integer k , we assign an operator on $C^\infty(Q)$ by

$$\xi_f^{(k)} = -\frac{i\hbar}{k} \xi_f.$$

Since $[\xi_f, X] = 0$, the operator $\xi_f^{(k)}$ restricts to an operator on each eigenspace \mathcal{E}_k of X having the form

$$\rho_k(f) = -\frac{i\hbar}{k} \overline{X}_f + m_f.$$

Evidently the map ρ_k satisfies the first Dirac axiom; moreover

$$\rho_k(\{f, g\}) = k [\rho_k(f), \rho_k(g)]_\hbar.$$

To verify that $\rho_k(f)$ is self-adjoint, it suffices to prove that $\rho_k(f)$ acts as a symmetric operator on the real subspace of \mathcal{E}_k . For this purpose, we note that if the hamiltonian vector field of

f is integrable (for example, if f is compactly supported), then the fact that ξ_f preserves φ implies that

$$\langle \xi_f \cdot u, v \rangle = -\langle u, \xi_f \cdot v \rangle$$

for real-valued functions u, v on Q . Combined with the definition of ρ_k , conjugate-symmetry of the inner-product implies that

$$\langle \rho_k(f) \cdot u, v \rangle = \langle u, \rho_k(f) \cdot v \rangle.$$

We now define the **prequantum line bundle** associated to Q as $E = E_{-1}$, equipped with a hermitian metric $\langle \cdot, \cdot \rangle$ compatible with its induced connection, and let \mathfrak{H}_P denote the L^2 completion of the space of compactly supported sections of E with respect to the inner-product

$$\langle s_1, s_2 \rangle = \int_P \langle s_1, s_2 \rangle \omega^n.$$

For each $f \in C^\infty(P)$, the operator $\rho_{-1}(f)$ extends to an essentially self-adjoint operator on \mathfrak{H}_P ; from the general remarks above, it follows that ρ_{-1} defines a prequantization of the Poisson algebra $C^\infty(P)$.

7.2 Polarizations and the metaplectic correction

The Kostant-Souriau prequantization of symplectic manifolds fails to satisfy the third Dirac axiom and is therefore not a quantization. In fact, the position and momentum operators

$$\rho(q) = -i\hbar \frac{\partial}{\partial p} + m_q \quad \rho(p) = -i\hbar \frac{\partial}{\partial q}$$

on $T^*\mathbb{R}$ both commute with $\partial/\partial p$ and therefore do not form a complete set. According to our earliest concepts of quantization (see the Introduction), the operator corresponding to q should act by multiplication alone, whereas \hat{q} involves the spurious $\partial/\partial p$ term. If we restrict to the p -independent subspace $C_q^\infty(\mathbb{R}^2) \simeq C^\infty(\mathbb{R})$, however, this difficulty is overcome: on $C_q^\infty(\mathbb{R}^2)$,

$$\rho(q) = m_q \quad \rho(p) = -i\hbar \frac{\partial}{\partial q}.$$

This association agrees with our earlier heuristic quantization of the classical position and momentum observables and suggests in general that the space of sections of a prequantum line bundle E is too “large” for the third Dirac axiom to be satisfied. Indeed, in quantum mechanics, wave functions depend on only half of the phase space variables, whereas the space of sections of E has the “size” of the space of functions on P . For an interpretation of this argument in terms of the Heisenberg uncertainty principle, see [53].

From the standpoint of WKB quantization, the appropriate quantum state space associated with a classical configuration space M is (the L^2 completion of) $|\Omega|_0^{1/2} M$, which we temporarily identify with $C^\infty(M)$ using a metric on M . Since the Liouville form α_M vanishes on each fiber of the projection $T^*M \xrightarrow{\pi} M$, we can interpret $C^\infty(M)$ as the space of sections of the prequantum line bundle E over T^*M which are parallel along each fiber of π . In other

words, the correct adjustment to the size of the prequantum Hilbert space is determined by the (canonical) foliation of T^*M by lagrangian submanifolds.

This adjustment is generalized in the framework of geometric quantization by the introduction of the following concept.

Definition 7.4 *A polarization of a symplectic manifold (P, ω) is an involutive lagrangian subbundle \mathcal{F} of $T_{\mathbb{C}}P$.*

Here, $T_{\mathbb{C}}P$ denotes the complexification of the tangent bundle of P , equipped with the complex-valued symplectic form $\omega_{\mathbb{C}}$ given by the complex-linear extension of ω to $T_{\mathbb{C}}P$. Integrability of \mathcal{F} means that locally on P there exist complex-valued functions whose hamiltonian vector fields (with respect to $\omega_{\mathbb{C}}$) span \mathcal{F} . The quantum state space associated to P is then given by the space of sections s of a prequantum line bundle E over P which are covariantly constant along \mathcal{F} . That is, for each complex vector field X on P lying in \mathcal{F} , we have $\nabla_X s = 0$.

Real polarizations

The standard polarization of a classical phase space T^*M is given by the complexification $\mathcal{F} = VM \oplus iVM$ of the vertical subbundle of $T(T^*M)$. In this case, the polarization \mathcal{F} satisfies $\overline{\mathcal{F}} = \mathcal{F}$, meaning that \mathcal{F}_p is a totally real subspace of $T_{\mathbb{C},p}(T^*M)$ for each $p \in T^*M$. In general, any polarization \mathcal{F} of a symplectic manifold P which satisfies $\overline{\mathcal{F}} = \mathcal{F}$ is the complexification of an integrable lagrangian subbundle of TP and is called a **real polarization** of P .

Example 7.5 The most basic examples of real polarizations are given by either the planes $q=\text{constant}$ or $p=\text{constant}$ in \mathbb{R}^{2n} . These correspond to the δ_q and $e^{i\pi\langle p, \cdot \rangle / \hbar}$ bases of $L^2(\mathbb{R}^n)$ respectively, in the sense that

$$f = \int_{\mathbb{R}^n} f(q) \delta(q) dq = (2\pi\hbar)^{-n/2} \int_{\mathbb{R}^n} e^{i\pi\langle p, q \rangle / \hbar} \hat{f} dp.$$

Covariant differentiation along vector fields X tangent to the $q=\text{constant}$ polarization \mathcal{F}_q is given by the ordinary Lie derivative with respect to X , and so the space of \mathcal{F}_q -parallel sections of the prequantum line bundle E identifies with the space of smooth functions on q -space. For the $p=\text{constant}$ polarization \mathcal{F}_p , we consider covariant derivatives of the form

$$\nabla_{\frac{\partial}{\partial q}} = \left(\frac{\partial}{\partial q} - p \frac{\partial}{\partial \sigma} \right).$$

A \mathcal{F}_p -parallel section ψ of E must therefore satisfy

$$\frac{\partial \psi}{\partial q} - 2\pi i p \psi$$

and is therefore of the form $\psi(q, p) = v(p)e^{i\pi\langle p, q \rangle / \hbar}$. Note that the \mathcal{F}_p representation of the free-particle $H(q, p) = p^2/2$ is $\rho(H) = m_{p^2/2}$, and so the p -polarization appears to be better adapted to this case.

△

A real polarization \mathcal{F} of a symplectic $2n$ -manifold P defines a subspace $\mathcal{F}(P) \subset C^\infty(P)$ of functions constant along the leaves of \mathcal{F} . According to the Hamilton-Jacobi theorem, the hamiltonian vector field of any member of $\mathcal{F}(P)$ is contained in \mathcal{F} . Thus $\{f, g\} = 0$ for any $f, g \in \mathcal{F}(P)$, i.e., $\mathcal{F}(P)$ is an abelian Poisson subalgebra of $C^\infty(P)$. By modifying the normal form results of Chapter 4, one can furthermore prove

Theorem 7.6 *Every real polarization is locally isomorphic to the q -constant polarization on \mathbb{R}^{2n} .*

By pulling back the position functions q_i to P , we have

Corollary 7.7 *For each $p \in P$, there exist $f_i \in \mathcal{F}(P)$, $i = 1, \dots, n$, such that the hamiltonian vector fields X_i span \mathcal{F} in a neighborhood of p .*

An **affine structure** on a manifold M is a curvature- and torsion-free connection on TM . Since any polarization-preserving symplectic transformation of $T^*\mathbb{R}^n \simeq \mathbb{R}^{2n}$ is affine on fibers (see Theorem 3.29), we have the following result.

Corollary 7.8 *Each leaf of any real polarization carries a natural affine structure.*

The main remark we wish to make for the moment is that our earlier prequantization ρ_{-1} of $C^\infty(P)$ now represents the subalgebra $\mathcal{F}(P) \subset C^\infty(P)$ as multiplication operators on $\mathfrak{H}_{\mathcal{F}}$. Since any $f \in \mathcal{F}(P)$ has the property that its hamiltonian vector field X_f lies completely within the polarization \mathcal{F} , it follows that $\nabla_{X_f}s = 0$ for all $s \in \mathfrak{H}_{\mathcal{F}}$. By the definition of ρ_{-1} , this implies that $\rho_{-1}(f) = m_f$, the multiplication operator on $\mathfrak{H}_{\mathcal{F}}$. More generally, those functions whose hamiltonian vector fields have flows which leave the polarization invariant (not necessarily leaf by leaf) are those which are “affine” along the leaves.

Example 7.9 Following [28], we call a polarization \mathcal{F} on P **fibrating** if each leaf of \mathcal{F} is simply-connected and the leaf space $P_{\mathcal{F}} = P/\mathcal{F}$ is a smooth manifold. In this case, the quotient map $p: P \rightarrow P_{\mathcal{F}}$ satisfies $\mathcal{F} = \ker p_*$, and so $p^*T^*P_{\mathcal{F}}$ identifies with the normal bundle $\mathcal{F}^\perp \subset T^*P$. Since \mathcal{F} is a lagrangian distribution, the map $\tilde{\omega}: TP \rightarrow T^*P$ sends \mathcal{F} to \mathcal{F}^\perp as well, thus defining a natural identification

$$\mathcal{F} \simeq p^*T^*P_{\mathcal{F}}.$$

A function $H: P \rightarrow \mathbb{R}$ constant on the leaves of \mathcal{F} induces a function $H_{\mathcal{F}}$ on $P_{\mathcal{F}}$ satisfying

$$dH = p^*dH_{\mathcal{F}}.$$

From this remark, it follows that a section s of \mathcal{F} over a leaf L is parallel with respect to the connection described above if and only if $(\tilde{\omega} \circ s)$ defines a fixed element of $T^*_{\{L\}}P_{\mathcal{F}}$.

Similarly, if E is a prequantum line bundle over P , we may use parallel sections of E over the leaves of \mathcal{F} to construct a foliation of the total space of E . The quotient of E by this

foliation defines a hermitian line bundle $E_{\mathcal{F}}$ over the leaf space $P_{\mathcal{F}}$ whose sections identify with \mathcal{F} -parallel sections of E .

The simplest example of a fibrating polarization is given by the vertical polarization \mathcal{F}_M of a cotangent bundle T^*M . The basic geometry of this situation remains the same if the symplectic structure on T^*M is perturbed by a form on M . More precisely, given a closed 2-form η on M , we can “twist” the standard symplectic structure on T^*M as follows:

$$\omega = \omega_M + \pi^*\eta.$$

Since $\pi^*\eta$ vanishes on the fibers of the projection, ω is a symplectic form on T^*M and $\ker \pi_*$ is again a polarization. This general type of symplectic manifold is known as a **twisted cotangent bundle**. A check of definitions shows that the standard and twisted symplectic structures on T^*M are equivalent precisely when their difference is the pull-back of an exact form on M . Moreover, the symplectic manifold $(T^*M, \omega_M + \pi^*\eta)$ is prequantizable if and only if the cohomology class of η in $H^2(M; \mathbb{R})$ is \hbar -integral. Finally, an application of Corollary 7.8 proves that if each leaf of a fibrating polarization \mathcal{F} on a symplectic manifold (P, ω) , is complete, then (P, ω) is symplectomorphic to a twisted symplectic structure on $T^*P_{\mathcal{F}}$.

△

Complex polarizations

Associated to any polarization \mathcal{F} are the distributions

$$D_{\mathbb{C}} = \mathcal{F} \cap \overline{\mathcal{F}} \quad E_{\mathbb{C}} = \mathcal{F} + \overline{\mathcal{F}}$$

which arise as the complexifications of distributions D, E in TP . Although their dimensions vary in general from point to point, D and E are pointwise ω -orthogonal, i.e.

$$D_x^{\perp} = E_x$$

for all $x \in P$. From the definition of \mathcal{F} , the distribution D is involutive. The polarization \mathcal{F} is called **strongly admissible** provided that E is involutive, the leaf spaces P_D and P_E are smooth manifolds, and the natural projection

$$P_D \xrightarrow{\pi} P_E$$

is a submersion. In this case, each fiber of π carries a Kähler structure, and geometric quantization attempts to construct a quantum state space for P from sections of an appropriate line bundle over P which are parallel along D and holomorphic along the fibers of π . For a discussion of quantization in this general setting, we refer to [53].

A polarization satisfying $\mathcal{F} \cap \overline{\mathcal{F}} = \{0\}$ is called **totally complex** and identifies with the graph of a complex structure J on P , i.e.,

$$\mathcal{F} = \{(v, iJv) : v \in TP\}$$

under the identification $T_{\mathbb{C}}P \simeq TP \oplus iTP$. We emphasize that J is a *genuine* complex structure on P due to the integrability condition on the polarization \mathcal{F} . Moreover, J is compatible in the usual sense with the symplectic form ω on P , so the hermitian metric $\langle \cdot, \cdot \rangle = g_J + i\omega$ is a Kähler structure on P . In this section, we will only cite two examples involving totally complex \mathcal{F} , in which case $P_D = P$ is Kähler, and $P_E = \{pt\}$.

Example 7.10 Consider the complex plane \mathbb{C} with its usual complex and symplectic structures. With respect to Darboux coordinates (q, p) , the Cauchy-Riemann operator is defined as

$$\frac{\partial}{\partial \bar{z}} = \frac{1}{2} \left(\frac{\partial}{\partial q} + i \frac{\partial}{\partial p} \right).$$

A function $f: \mathbb{C} \rightarrow \mathbb{C}$ is holomorphic if $\partial f / \partial \bar{z} = 0$.

If we identify sections of the (trivial) prequantum line bundle E with smooth, complex-valued functions on \mathbb{C} , then \mathcal{F} -parallel sections correspond to those functions annihilated by the covariant derivative

$$\nabla_{\frac{\partial}{\partial \bar{z}}} = 2 \frac{\partial}{\partial \bar{z}} + m_z.$$

To determine the general form of these sections, we note that $\nabla_{\frac{\partial}{\partial \bar{z}}} \psi = 0$ if and only if for some branch of the logarithm,

$$\frac{\partial}{\partial \bar{z}} \log \psi = -\frac{z}{2},$$

or

$$\log \psi = -\frac{z\bar{z}}{2} + h$$

for some holomorphic function h . Hence

$$\psi(z) = \varphi e^{-|z|^2/2}$$

with φ holomorphic is the general form for \mathcal{F} -parallel sections of E . The space $\mathfrak{H}_{\mathbb{C}}$ thus identifies with the space of holomorphic $\varphi: \mathbb{C} \rightarrow \mathbb{C}$ satisfying

$$\|\varphi\| = \int_{\mathbb{C}} |\varphi(z)|^2 e^{-|z|^2} dz < \infty,$$

or, in other words, the space of holomorphic functions which are square-integrable with respect to the measure $e^{-|z|^2} dz$, known as the **Fock** or **Bargman-Segal** space. For a function to be quantizable in this picture, its hamiltonian flow must preserve both the metric and symplectic structure of \mathbb{C} and therefore consist solely of euclidean motions. Among such functions are the usual position and momentum observables, as well as the harmonic oscillator.

△

Example 7.11 If (P, ω) is any prequantizable symplectic manifold with a totally complex polarization \mathcal{F} , then the prequantum line bundle E associated to ω can be given the structure of a holomorphic line bundle by taking the $(0, 1)$ component of the connection on E with

respect to the complex structure J arising from \mathcal{F} . The space $\mathcal{H}_{\mathcal{F}}$ of \mathcal{F} -parallel sections of E then equals the space of holomorphic sections of E and is therefore completely determined by the complex geometry of P .

Using some machinery from algebraic geometry (see [15]), it can be shown that for compact P , the space \mathcal{H}_P is finite-dimensional and that its dimension is given asymptotically by the symplectic volume of P . More precisely, we note that for $k \in \mathbb{Z}^+$, the line bundle $E^{\otimes k}$ defines a holomorphic prequantum line bundle associated with $(P, k\omega)$. For sufficiently large k , the dimension of the quantum state space \mathfrak{H}_{P_k} of holomorphic sections of $E^{\otimes k}$ is given by the Hirzebruch-Riemann-Roch formula:

$$\dim \mathfrak{H}_{P_k} = \int_P e^{k\omega} Td(P),$$

where Td denotes the **Todd polynomial** in the total Chern class of P . The first conclusion to be drawn from this fact is that for k large, $\dim \mathfrak{H}_{P_k}$ is a symplectic invariant of P independent of its complex structure. Roughly speaking, k plays the role of \hbar^{-1} , and so “large k ” means that we are approaching the classical limit. A second point to note is that $\dim \mathfrak{H}_{P_k}$ is a polynomial in k (or \hbar^{-1}) whose leading-order term is

$$\int_P \frac{(k\omega)^n}{n!} = k^n \text{vol } P.$$

Consequently the number of quantum states is determined asymptotically by the volume of P .

△

Metalinear structures and half-forms

For the remainder of this chapter, we will focus on the quantization of symplectic manifolds P equipped with a prequantum line bundle E and a “sufficiently nice” real polarization \mathcal{F} . Although the space $\mathfrak{H}_{\mathcal{F}}$ of \mathcal{F} -parallel sections of E appears to have the right “size” in the simplest examples, there are still several problems to be resolved before we have a suitable quantum state space. The first arises as soon as we attempt to define a pre-Hilbert space structure on $\mathfrak{H}_{\mathcal{F}}$. On P , the square of an \mathcal{F} -parallel section s is constant along the leaves of \mathcal{F} , and thus the integral

$$\|s\|^2 = \int_P \langle s, s \rangle \omega^n$$

diverges in general. On the other hand, there is no canonical measure on the leaf space $P_{\mathcal{F}}$ with which to integrate the induced function $\langle s, s \rangle$.

Regardless of how this first difficulty is resolved, we will again try to quantize the Poisson algebra $C^\infty(P)$ by representing elements of $\text{Aut}(P, \omega)$ as (projective) unitary operators on $\mathfrak{H}_{\mathcal{F}}$. The most obvious quantization of a symplectomorphism $f : P \rightarrow P$ (or more general canonical relation), however, is an operator from $\mathfrak{H}_{\mathcal{F}}$ to $\mathfrak{H}_{f(\mathcal{F})}$. If f does not preserve the polarization \mathcal{F} , then these spaces are distinct. Thus, we will need some means for canonically identifying the quantum state spaces $\mathfrak{H}_{\mathcal{F}}$ associated to different polarizations of P .

Finally, the polarization \mathcal{F} may have multiply-connected leaves, over most of which the prequantum line bundle may admit only the trivial parallel section. A preliminary solution to this problem is to admit “distributional” states, such as are given by \mathcal{F} -parallel sections of E whose restrictions to each leaf are smooth. Leaves on which E has trivial holonomy are then said to comprise the **Bohr-Sommerfeld subvariety** of the pair (P, \mathcal{F}) .

Example 7.12 Consider the punctured phase plane $\mathbb{R}^2 = \mathbb{R}^2 \setminus \{0\}$ polarized by level sets of the harmonic oscillator $H(q, p) = (q^2 + p^2)/2$. The holonomy of the usual prequantum line bundle on the level set $H^{-1}(E)$ is given by the mod- \mathbb{Z}_\hbar reduction of the area it encloses; thus the Bohr-Sommerfeld variety consists of those circles of energy $E = n\hbar$. As in the case of WKB prequantization, these energy levels do not correspond to actual physical measurements, and so we will again need to incorporate a sort of Maslov correction.

△

The solution to many of these difficulties relies on the use of a metaplectic structure on P , a concept which we now introduce. Let P be a principal G bundle over a manifold M . A **meta G -bundle** associated to P and a central extension

$$1 \rightarrow K \rightarrow \tilde{G} \xrightarrow{\rho} G \rightarrow 1$$

of G is a principal \tilde{G} bundle \tilde{P} over M together with a map $\Phi: \tilde{P} \rightarrow P$ satisfying the equivariance condition

$$\Phi(p \cdot a) = \Phi(p) \cdot \rho(a)$$

for all $a \in \tilde{G}$. Two meta G -bundles $(\tilde{P}_1, \Phi_1), (\tilde{P}_2, \Phi_2)$ over P are considered equivalent if there exists a \tilde{G} -equivariant diffeomorphism $\psi: \tilde{P}_1 \rightarrow \tilde{P}_2$ such that $\Phi_1 = \Phi_2 \circ \psi$.

Example 7.13 A riemannian structure together with an orientation of an n -manifold M defines a bundle of oriented orthonormal frames in TM , i.e. an $SO(n)$ -structure on M . A **spin structure** on M is then a meta $SO(n)$ -bundle corresponding to the extension

$$1 \rightarrow K \rightarrow Spin(n) \rightarrow SO(n) \rightarrow 1$$

given by the double cover $Spin(n)$ of $SO(n)$. An orientable manifold admits a spin structure if and only if its second Stiefel-Whitney class vanishes.

△

Example 7.14 Since $\pi_1(Sp(n)) \simeq \mathbb{Z}$, there exists a unique connected double-covering group of $Sp(n)$, known as the **metaplectic group** $Mp(n)$. A symplectic $2n$ -plane bundle F admits a reduction of its structure group from $GL(2n)$ to $Sp(n)$; a **metaplectic structure** on F then corresponds to a lifting of the symplectic frame bundle to a principal $Mp(n)$ bundle. In general, a symplectic vector bundle admits a metaplectic structure if and only if its second Stiefel-Whitney class vanishes.

A symplectic manifold P whose tangent bundle is equipped with a metaplectic structure is called a **metaplectic manifold**. If the Stiefel-Whitney class $w_2(P)$ is zero, metaplectic

structures are classified by the set $H^1(P; \mathbb{Z}_2)$ (see [28]). We emphasize that this classification depends on the bundle of metaplectic frames *and* its covering map to the bundle of symplectic frames. To see this explicitly in a special case, note that up to topological equivalence, the punctured plane \mathbb{R}^2 admits only trivial $Sp(1)$ and $Mp(1)$ principal bundles, since both $Sp(1)$ and $Mp(1)$ are connected. On the other hand, a metaplectic structure

$$\Phi: \mathbb{R}^2 \times Mp(1) \rightarrow \mathbb{R}^2 \times Sp(1)$$

is defined for any choice of continuous map $s: \mathbb{R}^2 \rightarrow Sp(1)$ by

$$\Phi(p, A) = (p, s(p) \cdot \rho(A))$$

for $A \in Mp(1)$. Two such maps s_0, s_1 define equivalent metaplectic structures if and only if $(s_1)^{-1} \cdot s_0$ admits a continuous lift $\tilde{s}: \mathbb{R}^2 \rightarrow Mp(1)$, in which case an equivalence is given by the map $\psi: \mathbb{R}^2 \times Mp(1) \rightarrow \mathbb{R}^2 \times Mp(1)$ defined as

$$\psi(p, A) = (p, \tilde{s}(p) \cdot A)$$

By the usual homotopy theory for continuous groups, this occurs precisely when $[s_0] = [s_1]$ as elements of $\pi_1(Sp(1))/\rho_{\#}\pi_1(Mp(1)) \simeq \mathbb{Z}_2$.

△

Example 7.15 Lying within the metaplectic group is a double-cover $ML(n)$ of the subgroup $Gl(n)$ preserving the usual lagrangian splitting of \mathbb{R}^{2n} . The double covering of $GL(n)$ then induced by the identification $GL(n) \simeq Gl(n)$ (see Example 3.1) is called the **metalinear group** $ML(n)$. It is trivial as a topological covering, but as a group it is not the direct product $GL(n) \times \mathbb{Z}_2$. More explicitly, the group $ML(n)$ is isomorphic to the direct product $GL^+(n) \times \mathbb{Z}_4$ with the covering map $ML(n) \xrightarrow{\rho} GL(n)$ given by

$$\rho(A, a) = A \cdot e^{i\pi a}.$$

A metalinear lifting of the frames of an n -plane bundle E is called a **metalinear structure** on E . If E is orientable, its structure group can be reduced to $GL^+(n)$, which can be interpreted as the identity component of $ML(n)$ in order to give a metalinear structure on E . More generally, a vector bundle admits a metalinear structure if and only if the square of its first Stiefel-Whitney class is zero (see [28]).

△

The importance of the metalinear group for our purposes is that it admits 1-dimensional representations which are not the lifts of representations of $GL(n)$. Bundles associated to metalinear structures via these representations will be the key to the metaplectic correction of the prequantization procedure. First note that by means of the quotient homomorphism

$$ML(n) \rightarrow ML(n)/GL^+(n) \simeq \mathbb{Z}_4,$$

we can associate a principal \mathbb{Z}_4 bundle to any metilinear structure $(MB(E), \Phi)$ on a vector bundle F . From the representation

$$a \mapsto e^{i\pi a/2}$$

of \mathbb{Z}_4 in $U(1)$, we then obtain a complex line bundle $\Lambda^{1/2}F$ called the **bundle of half-forms** associated to the triple $(F, MB(F), \Phi)$. The reason for this terminology is that the bundle $\Lambda^{1/2}F$ can be constructed directly from the metilinear frame bundle via the representation $ML(n) \rightarrow \mathbb{C}_*$ given by the square-root of $(\det \circ \rho)$

$$(A, a) \mapsto (\det A)^{1/2} e^{i\pi a/2}.$$

A section of $\Lambda^{1/2}F$ then identifies with a complex-valued function λ on $MB(F)$ satisfying

$$\lambda(e) = (\det A)^{1/2} e^{i\pi a/2} \cdot \lambda(e \cdot (A, a)),$$

and therefore represents, loosely speaking, the square-root of a volume form on F . By inverting and conjugating the preceding representation of $ML(n)$, one similarly defines the bundles $\Lambda^{1/2}\bar{F}$ and $\Lambda^{-1/2}F$ of conjugate and negative half-forms on F . Evidently, the product of two half-forms on F yields an n -form; similarly, a half-form λ and a conjugate half-form μ can be multiplied to give a 1-density $\lambda\mu$ on the bundle F .

An important link between metilinear and metaplectic structures on symplectic manifolds can be described as follows.

Theorem 7.16 *Let (P, ω) be a symplectic $2n$ -manifold and $L \subset TP$ a lagrangian subbundle. Then TP admits a metaplectic structure if and only if L admits a metilinear structure.*

Proof. If J is any ω -compatible almost complex structure on P , then $TP = L \oplus JL \simeq L \oplus L$. By the Whitney product theorem, we then have

$$w_2(P) = w_1(L)^2,$$

and our assertion follows from the remarks in the preceding examples.

More explicitly, we may use L and J to reduce the structure group of TP to the subgroup $Gl(n) \subset Sp(n)$ corresponding to $GL(n)$. The resulting frame bundle is isomorphic to the frame bundle of L , and thus, a metilinear structure on L can be enlarged to a metaplectic structure on TP , while a metaplectic structure on TP reduces to a metilinear structure on L .

□

Now let P be a metaplectic manifold equipped with a prequantum line bundle E and a real polarization \mathcal{F} . By the preceding theorem, \mathcal{F} inherits a metilinear structure from the metaplectic structure on TP , and thus $\Lambda^{-1/2}\mathcal{F}$ is defined; it is equipped with a natural flat connection inherited from the one on \mathcal{F} .

Definition 7.17 *The quantum state space $\mathfrak{H}_{\mathcal{F}}$ associated to P is the space of sections of $E \otimes \Lambda^{-1/2}\mathcal{F}$ which are covariantly constant and along each leaf of \mathcal{F} .*

In the following examples, we will sketch how in certain cases this definition enables us to overcome the difficulties mentioned at the beginning of this section.

Example 7.18 We first return to Example 7.12 to illustrate how the introduction of half-forms also enables us to incorporate the Bohr-Sommerfeld correction in the case of the 1-dimensional harmonic oscillator.

To quantize the punctured plane $\dot{\mathbb{R}}^2$ with its polarization by circles centered at the origin, we will use the trivial metaplectic structure on $\dot{\mathbb{R}}^2$ (see Example 7.14). In this case, both the trivial $Mp(1)$ and $Sp(1)$ bundles over $\dot{\mathbb{R}}^2$ can be identified with the set of triples $(r, \theta, a \cdot e^{i\phi})$, where $a, r > 0$ and $0 \leq \theta, \phi \leq 2\pi$, and the covering map $\dot{\mathbb{R}}^2 \times Mp(1) \rightarrow \dot{\mathbb{R}}^2 \times Sp(1)$ is given by

$$(r, \theta, a \cdot e^{i\phi}) \mapsto (r, \theta, a \cdot e^{2i\phi}).$$

The pull-back of $\dot{\mathbb{R}}^2 \times Sp(1)$ of $T\dot{\mathbb{R}}^2$ to the circle S_r of radius r can be reduced to the trivial $Gl(1)$ principal bundle over S_{r_0} by means of the lagrangian splitting $TS_{r_0} \oplus JTS_{r_0}$. In terms of the identifications above, the subbundle $\dot{\mathbb{R}}^2 \times Gl(1)$ equal the subset

$$\{(r, \theta, a \cdot e^{i\theta}) : r = r_0, a > 0, 0 \leq \theta \leq 2\pi\}$$

of $\dot{\mathbb{R}}^2 \times Sp(1)$ while its corresponding metalinear bundle equals the subset

$$\{(r, \theta, a \cdot e^{i(\theta/2+c\pi)}) : r = r_0, a > 0, 0 \leq \theta \leq 2\pi, c \in \{0, 1\}\}$$

of $\dot{\mathbb{R}}^2 \times Mp(1)$. Note that this is a *nontrivial* $Ml(1)$ bundle over S_{r_0} .

From the expression for the covering map given above, it follows easily that the parallel transport of a negative half-form associated to TS_{r_0} around S_{r_0} amounts to multiplication by -1 . Consequently, the bundle $E \otimes \Lambda^{-1/2}\mathcal{F}$ admits a parallel section over S_{r_0} precisely when

$$\pi r_0^2 = \int_{S_{r_0}} p dq = \pi \hbar (2n + 1)$$

for some integer n . That is, the Bohr-Sommerfeld variety consists of circles of energy $E = (n + 1/2)\pi\hbar$, in accordance with the corrected Bohr-Sommerfeld quantization conditions described in Chapter 4.

△

Example 7.19 Over each leaf L of the polarization \mathcal{F} , the Bott connection on TL described in Example 5.17 defines a foliation of the preimage of L in $B(\mathcal{F})$. As proven in [28], the quotient of $B(\mathcal{F})$ by this leaf-wise foliation defines a metalinear structure on the leaf space $P_{\mathcal{F}}$ in such a way that half-forms on $P_{\mathcal{F}}$ are in 1-1 correspondence with \mathcal{F} -parallel negative half-forms on \mathcal{F} . In this way, the quantum state space $\mathfrak{H}_{\mathcal{F}}$ identifies with the space of compactly supported sections of

$$E_{\mathcal{F}} \otimes \Lambda^{1/2}P_{\mathcal{F}}.$$

The latter space is equipped with a natural inner-product

$$\langle s_1 \otimes \lambda_1, s_2 \otimes \lambda_2 \rangle = \int_{P_{\mathcal{F}}} \langle s_1, s_2 \rangle \lambda_1 \bar{\lambda}_2.$$

Note that integration is well-defined, since $\lambda_1 \bar{\lambda}_2$ is a density on $P_{\mathcal{F}}$.

△

Blattner-Kostant-Sternberg kernels

If \mathcal{F}_1 and \mathcal{F}_2 are fibrating polarizations of P with the property that $\mathcal{F}_{1,x}$ is transverse to $\mathcal{F}_{2,x}$ at each $x \in P$, then the quantum state spaces $\mathfrak{H}_{\mathcal{F}_1}$ and $\mathfrak{H}_{\mathcal{F}_2}$ can be related as follows. By assumption, the symplectic form defines an isomorphism of \mathcal{F}_1 with \mathcal{F}_2^* ; similarly, the metaplectic structure on P defines a natural isomorphism between the space of half-forms on \mathcal{F}_2 and the space of conjugate half-forms on \mathcal{F}_1 (see [28]). Using this identification, we can pair a half-form λ_1 on \mathcal{F}_1 with a half-form λ_2 on \mathcal{F}_2 to obtain a function (λ_1, λ_2) on P . If we identify elements σ_1, σ_2 of $\mathfrak{H}_{\mathcal{F}_1}$ and $\mathfrak{H}_{\mathcal{F}_2}$ with \mathcal{F}_i -parallel sections $s_i \otimes \lambda_i$ of $E \otimes \Lambda^{1/2}\mathcal{F}_i$ on P , then a sesquilinear pairing $\mathfrak{H}_{\mathcal{F}_1} \times \mathfrak{H}_{\mathcal{F}_2} \rightarrow \mathbb{C}$ is defined by

$$\langle\langle \sigma_1, \sigma_2 \rangle\rangle = \frac{1}{(2\pi\hbar)^{n/2}} \int_P \langle s_1, s_2 \rangle (\lambda_1, \lambda_2) \omega^n,$$

where $\dim(P) = 2n$. This pairing is nondegenerate and thus determines a linear operator $B: \mathfrak{H}_{\mathcal{F}_2} \rightarrow \mathfrak{H}_{\mathcal{F}_1}$ satisfying

$$\langle\langle \sigma_1, \sigma_2 \rangle\rangle = \langle \sigma_1, B\sigma_2 \rangle$$

for all $\sigma_1 \in \mathfrak{H}_{\mathcal{F}_1}$ and $\sigma_2 \in \mathfrak{H}_{\mathcal{F}_2}$. This construction, which can be extended to more general pairs of polarizations, is due to Blattner, Kostant, and Sternberg defines a sesquilinear pairing on the associated quantum state spaces via a **Blattner-Kostant-Sternberg kernel**.

Example 7.20 The basic example of a Blattner-Kostant-Sternberg kernel arises from the classical Fourier transform. From Example 7.5, we recall that for the q -constant polarization \mathcal{F}_q of \mathbb{R}^{2n} , the space $\mathfrak{H}_{\mathcal{F}_q}$ consists of p -independent functions $\sigma_q(q, p) = u(q)$ on \mathbb{R}^{2n} , which we may identify with the space of smooth functions on q -space. Similarly, if \mathcal{F}_p is the p -constant polarization, then a smooth function $v(p)$ on p -space identifies with an element $\sigma_p(q, p) = v(p)e^{i\langle q, p \rangle/\hbar}$ of $\mathfrak{H}_{\mathcal{F}_p}$. The sesquilinear pairing described above is therefore given by

$$\langle\langle \sigma_q, \sigma_p \rangle\rangle = \frac{1}{(2\pi\hbar)^{n/2}} \int_{\mathbb{R}^{2n}} e^{i\langle q, p \rangle/\hbar} u(q) v(p) |dq dp|.$$

Note that this corresponds to the usual association of a distribution $\hat{v}(q)$ on q -space to a function $v(p)$ on p -space via the inverse asymptotic Fourier transform.

△

Given a symplectic manifold P with the structures above, geometric quantization attempts to represent the Poisson algebra $C^\infty(P)$ on $\mathfrak{H}_{\mathcal{F}}$ as follows. First, if the hamiltonian flow φ_t of $f \in C^\infty(P)$ preserves \mathcal{F} , then it lifts to a 1-parameter family $\tilde{\varphi}_t$ of operators on smooth sections of $E \otimes \Lambda^{1/2}\mathcal{F}$. The quantum operator associated to \mathcal{F} is then

$$\rho(f) \sigma = i\hbar \frac{d}{dt} (\tilde{\varphi}_t \sigma)|_{t=0}.$$

In particular, $\rho(f)$ acts as pointwise multiplication for any $f \in \mathcal{F}(P)$.

For those $f \in C^\infty(P)$ whose hamiltonian flow does not preserve \mathcal{F} , we assume that for $t \in (0, \varepsilon)$, the differential $D\varphi_t$ maps \mathcal{F} into a polarization \mathcal{F}_t for which there exists a Blattner-Kostant-Sternberg kernel and a corresponding operator $U_t: \mathfrak{H}_{\mathcal{F}_t} \rightarrow \mathfrak{H}_{\mathcal{F}}$. The flow φ_t defines an operator $\tilde{\varphi}_t: \mathfrak{H}_{\mathcal{F}} \rightarrow \mathfrak{H}_{\mathcal{F}_t}$, and the quantum operator $\rho(f)$ is defined as

$$\rho(f)\sigma = i\hbar \frac{d}{dt}((U_t \circ \tilde{\varphi}_t)\sigma)|_{t=0}.$$

We refer to [28] for further details.

7.3 Quantization of semi-classical states

In this brief section, we suppose that P is a metaplectic manifold with a prequantum line bundle E and a fibrating polarization \mathcal{F} . Our goal is to comment briefly on how an element of $\mathfrak{H}_{\mathcal{F}}$ can be constructed from a “semi-classical state” in P consisting of a lagrangian submanifold $L \subset P$ with some extra structure, such as a half-density or half-form with coefficients in a prequantum line bundle over P .

Example 7.21 If $L \subset P$ intersects each leaf of \mathcal{F} transversely in at most one point, then the section s of $E \otimes \Lambda^{-1/2}\mathcal{F}$ over L corresponding to a section of $E \otimes |\Lambda|^{1/2}L$ can be extended to a section \tilde{s} over P which is covariantly constant on each leaf and which vanishes on leaves which are disjoint from L . In this way, we can regard \tilde{s} as the quantization of the pair (L, s) .

We can generalize this viewpoint by considering a lagrangian submanifold L which has possibly multiple transverse intersections with the leaves of \mathcal{F} (that is, each intersection of L with a leaf of \mathcal{F} is transverse). If s is again a section of $E \otimes \Lambda^{-1/2}\mathcal{F}$ over L , then we can construct an element of $\mathfrak{H}_{\mathcal{F}}$ by “superposition”. More precisely, we think of L as the union of lagrangian submanifolds L_j , such that each L_j intersects any leaf of \mathcal{F} at most once. Applying the procedure above to each $(L_j, s|_{L_j})$ and summing the results produces an element of $\mathfrak{H}_{\mathcal{F}}$ as long as this sum converges.

Finally, if L is a union of leaves of \mathcal{F} and s is a section of $E \otimes \Lambda^{-1/2}\mathcal{F}$ over L which is covariantly constant along each leaf of \mathcal{F} , then we can pair it with arbitrary elements of $\mathfrak{H}_{\mathcal{F}}$ by integration over L , thus obtaining a linear functional on $\mathfrak{H}_{\mathcal{F}}$, which, via the inner product on $\mathfrak{H}_{\mathcal{F}}$, can be considered as a generalized section.

△

Example 7.22 To indicate briefly how “distributional” elements of $\mathfrak{H}_{\mathcal{F}}$ arising as in the preceding example are paired, we consider the following situation. Let (L, s) be a semi-classical state in \mathbb{R}^{2n} such that L equals the zero section of $\mathbb{R}^{2n} \simeq T^*\mathbb{R}^n$ and s is any section of $E \otimes \Lambda^{-1/2}\mathcal{F}_q$ over L . Similarly, let \tilde{L} be the fiber of $T^*\mathbb{R}^n$ over $0 \in \mathbb{R}^n$ along with a constant section \tilde{s} of $E \otimes \Lambda^{-1/2}\mathcal{F}_q$ over \tilde{L} . According to the preceding example, the quantization of (L, s) is given by an ordinary smooth function on \mathbb{R}^n , whereas the quantization of (\tilde{L}, \tilde{s}) is “concentrated” at the origin of \mathbb{R}^n .

The meaning of this last statement is made more precise by the fact that if $T: \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$ is a linear symplectomorphism, then the associated quantum operator $\mathfrak{H}_{\mathcal{F}_q} \rightarrow \mathfrak{H}_{T(\mathcal{F}_q)}$ is

required to be unitary. By choosing T so that both L and \tilde{L} are transverse to $T(\mathcal{F}_q)$, we can quantize (L, s) and (\tilde{L}, \tilde{s}) with respect to this new polarization to obtain smooth functions on the leaf space $\mathbb{R}_{T(\mathcal{F}_q)}^{2n}$ whose pairing is defined by integration. In terms of $\mathfrak{H}_{\mathcal{F}_q}$, this means that, up to a normalization, the quantum state obtained from (\tilde{L}, \tilde{s}) is a Dirac δ -function concentrated at the origin of \mathbb{R}^n .

△

8 Algebraic Quantization

An approach to quantization going back to Dirac and recently revived in [9][10][11] is based on the idea that the multiplicative structure of the $*$ -algebra of quantum observables is more central to quantization than the representation of observables as operators on Hilbert space, and that most of quantum mechanics can be done without regards for the precise nature of observables. Quantum observables comprise a noncommutative algebra \mathcal{A} which is assumed to belong to a family \mathcal{A}_\hbar of algebras, all with a common underlying vector space, but with an \hbar -dependent multiplication $*_\hbar$. As $\hbar \rightarrow 0$, the algebra \mathcal{A}_\hbar approaches a commutative algebra \mathcal{A}_0 , which may be interpreted as the algebra of functions on a classical phase space. Additionally, it is assumed that the $*_\hbar$ -commutator approaches the Poisson bracket in the limit of large quantum numbers:

$$\{f, g\} = \lim_{\hbar \rightarrow 0} (f *_\hbar g - g *_\hbar f) / i\hbar.$$

The abstract goal of algebraic quantization is to construct the family \mathcal{A}_\hbar of noncommutative algebras from a Poisson algebra. In this chapter, we sketch two approaches towards this goal, known as deformation quantization and the method of symplectic groupoids.

8.1 Poisson algebras and Poisson manifolds

The main objects in algebraic quantization theory are Poisson algebras and Poisson manifolds.

Definition 8.1 *A Poisson algebra is a real vector space \mathcal{A} equipped with a commutative, associative algebra structure*

$$(f, g) \mapsto fg$$

and a Lie algebra structure

$$(f, g) \mapsto \{f, g\}$$

which satisfy the compatibility condition

$$\{fg, h\} = f\{g, h\} + \{f, h\}g.$$

A Poisson manifold is a manifold P whose function space $C^\infty(P)$ is a Poisson algebra with respect to the usual pointwise multiplication of functions and a prescribed Lie algebra structure.

If P, Q are Poisson manifolds, a smooth map $\psi: P \rightarrow Q$ is called a **Poisson map** provided that it preserves Poisson brackets, i.e. $\{f, g\} = \{f \circ \psi, g \circ \psi\}$ for all $f, g \in C^\infty(Q)$. Similarly, ψ is called an **anti-Poisson map** if $\{f, g\} = -\{f \circ \psi, g \circ \psi\}$ for all $f, g \in C^\infty(Q)$.

On a Poisson manifold, the Leibniz identity implies that the Poisson bracket is given by a skew-symmetric contravariant tensor field π via the formula

$$\{f, g\} = \pi(df, dg).$$

The main examples of Poisson manifolds we will be concerned with are the following.

Example 8.2 1. Any smooth manifold M is a Poisson manifold when $C^\infty(M)$ is given the trivial bracket $\{f, g\} = 0$.

2. Any symplectic manifold is a Poisson manifold with respect to its standard Poisson structure

$$\{f, g\} = X_g \cdot f.$$

3. Let \mathfrak{g} be a finite-dimensional Lie algebra with dual \mathfrak{g}^* . The differential of a smooth function $F: \mathfrak{g}^* \times \mathfrak{g}^* \rightarrow \mathbb{R}$ is a map $DF: \mathfrak{g}^* \times \mathfrak{g}^* \rightarrow \mathfrak{g} \times \mathfrak{g}$ whose composition with the Lie bracket on \mathfrak{g} defines a smooth map $[DF]: \mathfrak{g}^* \times \mathfrak{g}^* \rightarrow \mathfrak{g}$. The **Lie-Poisson operator** on $C^\infty(\mathfrak{g}^* \times \mathfrak{g}^*)$ is the differential operator \mathcal{D} defined by

$$\mathcal{D}F = \frac{1}{2} \langle [DF]_{(x,y)}, x + y \rangle.$$

The dual space \mathfrak{g}^* is then a Poisson manifold when equipped with the **Lie-Poisson bracket**

$$\{f, g\} = \Delta^* \mathcal{D}(f \boxtimes g),$$

where $\Delta: \mathfrak{g}^* \rightarrow \mathfrak{g}^* \times \mathfrak{g}^*$ is the diagonal and $f, g \in C^\infty(\mathfrak{g}^*)$. More concretely,

$$\{f, g\}(\mu) = \mu([Df, Dg]),$$

where $\mu \in \mathfrak{g}^*$, and $Df, Dg: \mathfrak{g}^* \rightarrow \mathfrak{g}^{**} \simeq \mathfrak{g}$ are the differentials of f and g .

4. If V is a finite-dimensional vector space and π is a skew-symmetric bilinear form on V^* , then

$$\{f, g\} \stackrel{def}{=} \pi(df, dg)$$

defines a Poisson algebra structure on $C^\infty(V)$, making V a Poisson manifold.

△

8.2 Deformation quantization

The aim of deformation quantization is to describe the family $*_{\hbar}$ of quantum products on a Poisson algebra \mathcal{A} as an asymptotic series (in \hbar) of products on \mathcal{A} . In accordance with the introductory remarks above, the zeroth and first order terms of this series are determined by the Poisson algebra structure of \mathcal{A} ; the role of the higher-order terms is roughly speaking to give a more precise \hbar -dependent path from classical to quantum mechanics.

Definition 8.3 *Let \mathcal{A} be a complex vector space equipped with a commutative associative algebra structure, and let $*_{\hbar}$ be a family of associative multiplications on \mathcal{A} given by a formal power series*

$$f *_{\hbar} g = \sum_{j=0}^{\infty} B_j(f, g) \hbar^j$$

where each $B_j: \mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A}$ is a bilinear map. Then $*_{\hbar}$ is called a ***-deformation** of \mathcal{A} if

1. $B_0(f, g)$ equals the product in \mathcal{A} .
2. $B_j(g, f) = (-1)^j B_j(f, g)$
3. $B_j(1, f) = 0$ for $j \geq 1$
4. B_j is a differential operator in each argument.
5. $(f *_\hbar g) *_\hbar h = f *_\hbar (g *_\hbar h)$

For condition (5) to make sense, we must extend the product $*_\hbar$ in the obvious way from \mathcal{A} to the space $\mathcal{A}[[\hbar]]$ of formal power series. We emphasize that the power series above is not in general assumed to converge for $\hbar \neq 0$. Instead it should be viewed as an ‘‘asymptotic expansion’’ for the product, and manipulations of this series will be purely formal.

In any case, conditions (1-5) imply that

$$\{f, g\} \stackrel{def}{=} \frac{1}{2i} B_1(f, g)$$

defines a Poisson algebra structure on \mathcal{A} . This observation suggests the question of whether every Poisson structure on a function algebra $C^\infty(P)$ can be realized as the first order term in a $*$ -deformation of \mathcal{A} .

Example 8.4 If V is a finite-dimensional real vector space and $\pi : V^* \times V^* \rightarrow \mathbb{R}$ is a skew-symmetric bilinear form on V^* , then the Hessian of π at $0 \in V^* \times V^*$ is a linear map $A = \mathcal{H}\pi : V^* \times V^* \rightarrow V \times V$. We define the **Poisson operator** associated to π as the second-order differential operator on $C^\infty(V \times V)$

$$\mathcal{D}_\pi = \langle A(\partial/\partial y^*, \partial/\partial z^*), (\partial/\partial y, \partial/\partial z) \rangle,$$

where (y, z) are linear coordinates on $V \times V$ arising from a single set of linear coordinates on V , and (y^*, z^*) are dual to (y, z) . The **Moyal-Weyl operator** is then the pseudodifferential operator given by exponentiation:

$$\mathcal{M}_{\pi, \hbar} = e^{-i\hbar \mathcal{D}_\pi/2}.$$

Pointwise multiplication of functions in $f, g \in C^\infty(\mathbb{R}^{2n})$ can be defined as the pull-back

$$f \cdot g = \Delta^* f \boxtimes g,$$

where $\Delta : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n} \times \mathbb{R}^{2n}$ is the diagonal embedding, and $(f \boxtimes g)(y, z) = f(y)g(z)$ for $f, g \in C^\infty(\mathbb{R}^{2n})$. Similarly, a straightforward computation shows that the Poisson operator is related to the Poisson bracket induced by π via the equation

$$\{f, g\} = \Delta^* \mathcal{D}(f \boxtimes g).$$

The diagonal pull-back

$$f *_\hbar g = \Delta^* \mathcal{M}_\hbar(f \boxtimes g)$$

of the Moyal operator defines a $*$ -deformation of the Poisson algebra $C^\infty(V)$ called the **Moyal-Weyl product**. In terms of the linear coordinates (y, z) on $V \times V$, the operator B_j in the expansion of the Moyal-Weyl product is

$$B_j(f, g) = \frac{1}{j!} \left(\frac{i}{2} \sum_{r,s} \pi_{r,s} \frac{\partial}{\partial y_r} \frac{\partial}{\partial z_s} \right)^j f(y)g(z) \Big|_{y=z=x}.$$

If π is nondegenerate, then an application of the principle of stationary phase shows that an integral expression for $f *_{\hbar} g$ is given by

$$(f *_{\hbar} g)(x) = \int \int e^{i\hbar Q(y-x, z-x)} f(y) g(z) dy dz,$$

where Q is the skew-symmetric bilinear form on V induced by π .

△

A Poisson manifold P is called **regular** if its Poisson tensor π has constant rank, in which case a theorem of Lie [39] asserts that P is locally isomorphic to a vector space with a constant Poisson structure. In view of Example 8.4, any regular Poisson manifold is *locally* deformation quantizable. To construct a $*$ -product on all of $C^\infty(P)$, one may therefore attempt to “patch together” the local deformations to arrive at a global $*$ -product.

This technique has succeeded. A theorem of DeWilde and Lecomte [19] asserts that the Poisson algebra of any finite-dimensional symplectic manifold admits a $*$ -deformation. Using similar techniques, Mélotte [44] extended their result to arbitrary regular Poisson manifolds. A simplified proof of these results has recently been given by Fedosov [23] (see [68] for a survey of deformation quantization, emphasizing Fedosov’s construction).

8.3 Symplectic groupoids

The method of symplectic groupoids also attempts to directly construct a noncommutative algebra \mathcal{A}_{\hbar} of quantum observables without explicitly identifying a quantum state space. Unlike deformation quantization, however, this approach involves a geometric procedure which attempts to construct \mathcal{A}_{\hbar} for a particular value of \hbar and in particular incorporates geometric objects with certain quantum properties.

Groupoids

In this section we collect some basic definitions and examples of groupoids and their counterparts in symplectic geometry.

Definition 8.5 *A groupoid is a set Γ endowed with a product map $(x, y) \rightarrow xy$ defined on a subset $\Gamma_2 \subset \Gamma \times \Gamma$ called the set of composable pairs, and an inverse map $\iota: \Gamma \rightarrow \Gamma$ satisfying the conditions*

1. $\iota^2 = \text{id}$

2. If $(x, y), (y, z) \in \Gamma_2$, then $(xy, z), (x, yz) \in \Gamma_2$, and $(xy)z = x(yz)$.
3. $(\iota(x), x) \in \Gamma_2$ for all $x \in \Gamma$, and if $(x, y) \in \Gamma_2$, then $\iota(x)(xy) = y$.
4. $(x, \iota(x)) \in \Gamma_2$ for all $x \in \Gamma$, and if $(z, x) \in \Gamma_2$, then $(zx)\iota(x) = z$.

Note that by (3), the map ι is bijective and thus inverses in Γ are unique. An element x of Γ can be thought of as an arrow with source $\alpha(x) = \iota(x)x$ and target $\beta(x) = x\iota(x)$; a pair (x, y) then belongs to Γ_2 if and only if the source of y equals the target of x . The set Γ_0 of all sources (and targets) is called the **base** of Γ , and Γ is said to be a groupoid over Γ_0 . Elements of Γ_0 are units of Γ in the sense that $x\alpha(x) = x$ and $\beta(x)x = x$ for all $x \in \Gamma$. Finally, the **multiplication relation** of Γ is the subset

$$\mathfrak{m} = \{(xy, x, y) : (x, y) \in \Gamma_2\}$$

of $\Gamma \times \Gamma \times \Gamma$. In abstract terms, a groupoid is a small category in which all morphisms have inverses.

Example 8.6 1. Any group is a groupoid over its identity element, and conversely, any groupoid whose base is a singleton comprises a group.

2. A disjoint union of groupoids is a groupoid over the union of their bases. If Γ is a groupoid with base Γ_0 and $\Gamma'_0 \subset \Gamma_0$, then $\Gamma' = \{x \in \Gamma : \alpha(x), \beta(x) \in \Gamma'_0\}$ is a groupoid over Γ'_0 .

3. Combining (1) and (2), we see that any vector bundle E defines a groupoid $\Gamma(E)$ over its zero section.

4. The **pair groupoid** associated to a set X consists of $\Gamma = X \times X$, endowed with the multiplication $(x, y) \cdot (y, z) = (x, z)$. Thus Γ_0 is the diagonal, and α, β are the projections $\alpha(x, y) = (y, y)$ and $\beta(x, y) = (x, x)$. In this groupoid, there is exactly one arrow from any object to another.

△

We will be interested in groupoids with some geometric structures. Maintaining the notation above, we make the following definition.

Definition 8.7 A groupoid Γ is called a **Lie groupoid** if

1. Γ_0 is a submanifold of Γ .
2. The mappings $\alpha, \beta: \Gamma \rightarrow \Gamma_0$ are submersions.
3. Multiplication $\Gamma_2 \rightarrow \Gamma$ and inversion $\Gamma \xrightarrow{\iota} \Gamma$ are smooth.

Condition (2) implies that the map $\alpha \times \beta$ is transverse to the diagonal Δ in $\Gamma_0 \times \Gamma_0$, and so both $\Gamma_2 = (\alpha \times \beta)^{-1}(\Delta) \subset \Gamma \times \Gamma$ and the multiplication relation $\mathfrak{m} \subset \Gamma \times \Gamma \times \Gamma$ are smooth submanifolds.

A submanifold L of Γ is called **unitary** if the restriction of α and β to L are diffeomorphisms $L \rightarrow \Gamma_0$. The unitary submanifolds form a group under the natural multiplication of subsets.

Definition 8.8 A Lie groupoid Γ is called a **symplectic groupoid** if Γ is a symplectic manifold and the multiplication relation \mathbf{m} is a lagrangian submanifold of $\Gamma \times \bar{\Gamma} \times \bar{\Gamma}$.

Two immediate consequences of this definition and the calculus of canonical relations are that Γ_0 is lagrangian, and $\iota: \Gamma \rightarrow \Gamma$ is anti-symplectic, i.e., its graph is a lagrangian submanifold of $\Gamma \times \Gamma$. Thus, a symplectic groupoid Γ is characterized by the three canonical relations (recall that Z is a point):

$$\Gamma_0 \in \text{Hom}(Z, \Gamma) \quad L_\iota \in \text{Hom}(\bar{\Gamma}, \Gamma) \quad \mathbf{m} \in \text{Hom}(\Gamma \times \Gamma, \Gamma)$$

linked by the equation

$$\Gamma_0 = \mathbf{m} \circ L_\iota.$$

As a consequence of the axioms and the assumption that \mathbf{m} is lagrangian, there is a unique Poisson structure on the base Γ_0 of a symplectic groupoid Γ such that $\alpha: \Gamma \rightarrow \Gamma_0$ is a Poisson map and $\beta: \Gamma \rightarrow \Gamma_0$ is anti-Poisson. This Poisson structure gives meaning to the following concept.

Definition 8.9 A symplectic groupoid Γ is said to **integrate** a Poisson manifold P if there exists a Poisson isomorphism from the base Γ_0 of Γ onto P .

If there exists a groupoid Γ which integrates P , we will say that P is **integrable** and refer to Γ as a symplectic groupoid *over* P .

Example 8.10 1. A simple example of a Lie groupoid is given by any Lie group. From the requirement that the base of a symplectic groupoid be a lagrangian submanifold, only discrete Lie groups can be symplectic groupoids.

2. If M is any manifold with the zero Poisson structure, then T^*M , equipped with the groupoid structure of a vector bundle and its standard symplectic structure, defines a symplectic groupoid over M .

3. If P is any symplectic manifold, then the pair groupoid structure on $P \times \bar{P}$ defines a symplectic groupoid over P . Its unitary lagrangian submanifolds are precisely the graphs of symplectomorphisms of P .

4. If \mathfrak{g} is any Lie algebra, then the dual space \mathfrak{g}^* with its Lie-Poisson structure is integrable. For any Lie group G whose Lie algebra is \mathfrak{g} , we define a groupoid structure on T^*G by taking α and β to be the right and left translations of covectors to the fiber at the identity $e \in G$. A simple computation shows that the base of T^*G is its fiber at the identity, while L_ι and \mathbf{m} identify with the conormal bundles to the inversion and multiplication relations of G under the identifications $T^*G \times T^*G \simeq T^*(G \times G)$ $T^*G \times \bar{T^*G} \times \bar{T^*G} \simeq T^*(G \times G \times G)$. Equipped with these operations and its usual symplectic structure, T^*G is a symplectic groupoid over \mathfrak{g}^* .

△

Quantization via symplectic groupoids

From our discussion of geometric quantization, we know that certain symplectic manifolds quantize to give vector spaces V , and lagrangian submanifolds correspond to elements in V . If we wish V to be an associative $*$ -algebra with unit element, like the algebras of quantum mechanics, then the underlying symplectic manifold must possess a groupoid structure compatible with its symplectic structure.

The first step in the quantization of a Poisson manifold P by the method of symplectic groupoids is to construct a symplectic groupoid $(\Gamma, \Gamma_0, \iota, \mathfrak{m})$ over P . This is known as the integration problem for Poisson manifolds. Using the techniques of geometric quantization (prequantizations, polarizations), we then attempt to associate a vector space \mathcal{A}_Γ to Γ in such a way that the canonical relations Γ_0, L_ι , and \mathfrak{m} quantize as elements of $\mathcal{A}_\Gamma, \mathcal{A}_\Gamma \otimes \mathcal{A}_\Gamma^*$, and $\mathcal{A}_\Gamma \otimes \mathcal{A}_\Gamma^* \otimes \mathcal{A}_\Gamma^*$ which define the structure of an associative $*$ -algebra on \mathcal{A}_Γ . If successful, it is in this sense that the symplectic groupoid Γ represents a classical model for the quantum algebra \mathcal{A}_Γ .

To conclude this chapter, we will give several examples illustrating the spirit of the symplectic groupoid method. Throughout, we will only deal with groupoids Γ equipped with reasonable (e.g. fibrating) polarizations, so that the construction of the vector space \mathcal{A}_Γ follows unambiguously from the procedure defined in the preceding chapter.

Example 8.11 Consider a trivial Poisson manifold M and its associated symplectic groupoid T^*M . As described in the preceding example, the identity relation Γ_0 coincides with the zero section of T^*M ; under the identifications $T^*M \times T^*M \simeq T^*(M \times M)$ and $T^*M \times \overline{T^*M} \times \overline{T^*M} \simeq T^*(M \times M \times M)$, the relations L_ι and \mathfrak{m} identify with the conormal bundles of the diagonals $\Delta_2 \subset M \times M$ and $\Delta_3 \subset M \times M \times M$ respectively.

As in Chapter 7, we may identify the quantum Hilbert spaces associated to $T^*M, T^*(M \times M)$, and $T^*(M \times M \times M)$ with (completions of) the function spaces $C^\infty(M), C^\infty(M \times M)$ and $C^\infty(M \times M \times M)$. The relation Γ_0 then quantizes as the function 1 on M . Since L_ι and \mathfrak{m} are the conormal bundles of the diagonals in $M \times M$ and $M \times M \times M$, respectively, our heuristic discussion in Section 7.3 shows that, after an appropriate normalization, these relations are quantized by the δ -functions $\delta(x, y)$ and $\delta(x, z)\delta(y, z)$ supported on the diagonals Δ_2 and Δ_3 respectively. Thus, the quantization of the groupoid Γ yields the usual identity element, complex conjugation, and pointwise multiplication in the associative $*$ -algebra $C^\infty(M, \mathbb{C})$.

△

Example 8.12 Arguing as in the preceding example, we find that the quantization of the canonical relations Γ_0, L_ι , and \mathfrak{m} associated to the groupoid T^*G of Example yields the distributions $\delta(e)$ on G , $\delta(g_1, g_1^{-1})$ on $G \times G$, and $\delta(g_1 g_2, g_1, g_2)$ on $G \times G \times G$. If the Haar measure on G is used to identify the quantum Hilbert space \mathfrak{H}_G with $C^\infty(G, \mathbb{C})$, then the relations Γ_0, L_ι and \mathfrak{m} quantize as evaluation (at $e \in G$), anti-involution $f(g) \mapsto \overline{f(g^{-1})}$, and convolution.

There is actually a flaw in the preceding two examples, since geometric quantization produces half-densities rather than functions, and furthermore, the natural domain of the

convolution operation on a group consists not of functions but of densities. It appears then that the construction of V itself, and not just the multiplication, should depend on the groupoid structure on Γ .

Although the two multiplications associated with the two groupoid structures on T^*G described above live on different spaces, it is possible to relate them more closely by dualizing one of them, say convolution. In this way, we obtain the *coproduct* Δ on $C^\infty(G)$ defined as a map from $C^\infty(G)$ to $C^\infty(G) \otimes C^\infty(G) = C^\infty(G \times G)$ by the formula $(\Delta f)(x, y) = f(xy)$. The coproduct satisfies a coassociative law and is related to pointwise multiplication by the simple identity $\Delta(fg) = \Delta(f)\Delta(g)$, making $C^\infty(G)$ into a **Hopf algebra**. The compatibility between the two structures on $C^\infty(G)$ reflects a compatibility between the two groupoid structures on T^*G .

△

Example 8.13 If we take $P = \mathbb{R}^{2n}$ with its standard symplectic structure, there is a polarization of $P \times \overline{P}$ which does not depend on any polarization of P but only on the affine structure. In fact, the polarization comes from the isomorphism of $\Gamma = \mathbb{R}^{2n} \times \overline{\mathbb{R}^{2n}}$ with $T^*\mathbb{R}^{2n}$ given by

$$(x, y) \mapsto ((x + y)/2, \tilde{\omega}_n(y - x)).$$

Using this polarization to quantize Γ , we get as V the space of smooth functions on the diagonal, which we identify with \mathbb{R}^{2n} itself, and the multiplication on V turns out to be the Moyal product (see Example 8.4).

△

Example 8.14 If V is any finite-dimensional vector space, then a skew-adjoint linear map $\pi : V^* \rightarrow V$ defines a skew-symmetric bilinear form on V^* and thus a translation-invariant Poisson structure on V . If \mathbb{T} is a torus equal to the quotient of V by some lattice, then \mathbb{T} inherits a translation-invariant Poisson structure from V . A symplectic groupoid which integrates \mathbb{T} is given by the cotangent bundle $T^*\mathbb{T}$, with Γ_0 equals the zero section of $T^*\mathbb{T}$, and L_ι equals the conormal bundle of the diagonal in $\mathbb{T} \times \mathbb{T}$. To describe the multiplication relation, we identify $T^*\mathbb{T}$ with $\mathbb{T} \times V^*$ and let $\mathfrak{m} \subset T^*(T \times T \times T)$ consist of all triples (q, q', q'', p, p', p'') such that $q'' = q' + \frac{1}{2}T(p' + p'')$ and

$$q = q' + \frac{1}{2}T(p'') \qquad p = p' + p''.$$

When the map T is zero, the Poisson structure on \mathbb{T} is trivial, and the groupoid product quantized to the usual pointwise multiplication of functions on \mathbb{T} . Otherwise one gets a noncommutative multiplication on $C^\infty(\mathbb{T})$ which is precisely that of (the functions on) a **noncommutative torus**, one of the basic examples of noncommutative geometry. (The relation between Poisson tori and noncommutative tori was studied from the point of view of deformation quantization in [50]).

△

A Densities

An n -form ν on an n -dimensional manifold M can be viewed as a scalar function on the space of bases in the tangent bundle which satisfies

$$\nu(eA) = \nu(e) \cdot \det(A),$$

where $e = (e_1, \dots, e_n)$ is any frame in a tangent space of M and $A = (a_{ij})$ is any invertible $n \times n$ matrix. Because the “change of variables” formula for integration involves absolute values of Jacobians, integration of n -forms on M requires a choice of orientation. The use of densities instead of forms circumvents this need.

A *density* on a real vector space V of dimension n is a complex-valued function η , defined on the set $\mathcal{B}(V)$ of bases in V , which satisfies $\eta(eA) = \eta(e) \cdot |\det(A)|$. The collection of such functions is denoted $|\Lambda|V$. This concept can be generalized as follows.

Definition A.1 For $\alpha \in \mathbb{C}$, an α -density on V is a map $\lambda : \mathcal{B}(V) \rightarrow \mathbb{C}$ such that

$$\lambda(eA) = \lambda(e) \cdot |\det(A)|^\alpha.$$

We denote the vector space of α -densities on V by $|\Lambda|^\alpha V$. Since $GL(V)$ acts transitively on $\mathcal{B}(V)$, an α -density is determined by its value on a single basis. As a result, $|\Lambda|^\alpha V$ is a 1-dimensional complex vector space.

Operations on densities

1. If $\sigma \in |\Lambda|^\alpha V$ and $\beta \in \mathbb{C}$, then σ^β is a well-defined $\alpha\beta$ -density on V .
2. A linear map $T : V \rightarrow V^*$ induces a real-valued 1-density $\|T\|$ on V given by

$$\|T\| e \stackrel{\text{def}}{=} |\det(\langle Te_i, e_j \rangle)|^{1/2}.$$

Equivalently, any real bilinear form ω on V induces a 1-density $\|\tilde{\omega}\|$ on V .

3. Multiplication of densities is defined by multiplication of their values and gives rise to a bilinear map:

$$|\Lambda|^\alpha V \times |\Lambda|^\beta V \rightarrow |\Lambda|^{\alpha+\beta} V.$$

4. If W is a subspace of V , then a basis of V , unique up to transformation by a matrix of determinant ± 1 , is determined by a choice of bases for W and V/W . Consequently, there is a natural product

$$|\Lambda|^\alpha W \times |\Lambda|^\alpha(V/W) \rightarrow |\Lambda|^\alpha V$$

In particular, if $V = V_1 \oplus V_2$, then there is a natural product

$$|\Lambda|^\alpha V_1 \times |\Lambda|^\alpha V_2 \rightarrow |\Lambda|^\alpha V.$$

5. Operations 3 and 4 induce natural isomorphisms

$$|\Lambda|^\alpha V \otimes |\Lambda|^\beta V \rightarrow |\Lambda|^{\alpha+\beta} V$$

$$|\Lambda|^\alpha W \otimes |\Lambda|^\alpha (V/W) \rightarrow |\Lambda|^\alpha V.$$

$$|\Lambda|^\alpha V_1 \otimes |\Lambda|^\alpha V_2 \rightarrow |\Lambda|^\alpha (V_1 \oplus V_2)$$

6. If $T : V \rightarrow V'$ is an isomorphism, there is a well-defined isomorphism

$$T_* : |\Lambda|^\alpha V \rightarrow |\Lambda|^\alpha V'.$$

7. A natural map $\mathcal{B}(V) \xrightarrow{*} \mathcal{B}(V^*)$ is defined by associating to each basis e of V its dual basis e^* . Since $(eA)^* = e^*(A^*)^{-1}$ for any $A \in GL(V)$, the map $*$ gives rise to a natural isomorphism

$$|\Lambda|^\alpha V \rightarrow |\Lambda|^{-\alpha} V^*.$$

Example A.2 Suppose that A, C are vector spaces and

$$0 \rightarrow A \rightarrow A \oplus C \xrightarrow{F} C^* \rightarrow 0$$

is an exact sequence such that $F|_C = T$. By the operations described above, we obtain an isomorphism

$$|\Lambda|^\alpha A \simeq |\Lambda|^\alpha A \otimes |\Lambda|^{2\alpha} C$$

given explicitly by

$$\sigma \mapsto \sigma \otimes \|T\|^{2\alpha} = |\det_\theta T|^\alpha \sigma \otimes \theta^{2\alpha},$$

where θ is any real-valued 1-density on C , and the positive real number $|\det_\theta T|$ is defined by the equation

$$|\det_\theta T|^{1/2} \theta = \|T\|.$$

△

It is easy to check that the association $V \mapsto |\Lambda|^\alpha V$ defines a differentiable functor, so we can associate the α -density bundle $|\Lambda|^\alpha E$ to any smooth vector bundle E over a manifold M . The remarks above imply that if

$$0 \rightarrow E \rightarrow F \rightarrow G \rightarrow 0$$

is an exact sequence of vector bundles over M , then there is a natural density-bundle isomorphism

$$|\Lambda|^\alpha E \otimes |\Lambda|^\alpha G \simeq |\Lambda|^\alpha F.$$

We denote by $|\Omega|^\alpha E$ the vector space of smooth sections of $|\Lambda|^\alpha E$ and by $|\Omega|_c^\alpha E$ the space of smooth, compactly-supported sections. The density spaces associated to the tangent bundle of M are denoted $|\Omega|^\alpha M$. An n -form on an n -manifold M induces an α -density $|\nu|^\alpha$ in the manner of (4) above.

A natural mapping $|\Omega|_c^1 M \rightarrow \mathbb{C}$ is defined by integration

$$\sigma \mapsto \int_M \sigma.$$

Similarly, a pre-Hilbert space structure on the space $|\Omega|_c^{1/2} M$ of smooth, compactly-supported half-densities on M is defined by

$$\langle \sigma, \tau \rangle = \int_M \sigma \bar{\tau}.$$

The completion \mathfrak{H}_M of this space is called the **intrinsic Hilbert space** of M .

B The method of stationary phase

We begin with a few useful formulas involving the asymptotic Fourier transform. Let \mathcal{S} denote the usual Schwartz space of rapidly decreasing, complex valued functions on \mathbb{R}^n . For $u \in \mathcal{S}$, the asymptotic Fourier transform and its inverse are defined respectively as

$$\begin{aligned} (\mathcal{F}_\hbar u)(\xi) &= (2\pi\hbar)^{-\frac{n}{2}} \int_{\mathbb{R}^n} e^{-i\langle x, \xi \rangle / \hbar} u(x) dx \\ (\mathcal{F}_\hbar^{-1} v)(x) &= (2\pi\hbar)^{-\frac{n}{2}} \int_{(\mathbb{R}^n)^*} e^{i\langle x, \xi \rangle / \hbar} v(\xi) d\xi \end{aligned}$$

To see that these transforms are actually inverse to one another, note that the change of variables $\eta = \xi / \hbar$ gives

$$(2\pi\hbar)^{-\frac{n}{2}} \int_{(\mathbb{R}^n)^*} e^{i\langle x, \xi \rangle / \hbar} (\mathcal{F}_\hbar u)(\xi) d\xi = (2\pi)^{-n} \int \int e^{i\langle x-y, \eta \rangle} u(y) dy d\eta,$$

which equals $u(x)$ by the usual Fourier inversion formula. This observation also verifies the asymptotic inversion formula:

$$u(x) = (2\pi\hbar)^{-n} \int \int e^{i\langle x-y, \xi \rangle / \hbar} u(y) dy d\xi.$$

A simple application of this formula shows that the asymptotic differential operator $D_j = -i\hbar\partial_j$ satisfies the familiar equations

$$\mathcal{F}_\hbar(D_j u) = \xi_j \mathcal{F}_\hbar u \quad \mathcal{F}_\hbar(x_j u) = -D_j \mathcal{F}_\hbar u.$$

A similar check of definitions proves the asymptotic Parseval formula:

$$\int_{\mathbb{R}^n} u \bar{v} dx = \int_{(\mathbb{R}^n)^*} \mathcal{F}_\hbar u \overline{\mathcal{F}_\hbar v} d\xi.$$

We study next the asymptotic behavior of integrals of the form

$$I_\hbar = \int_{\mathbb{R}^n} e^{iR(x)/\hbar} a(x) |dx| \quad a \in C_0^\infty(\mathbb{R}^n), R \in C^\infty(\mathbb{R}^n)$$

as $\hbar \rightarrow 0$. As a first step, we will prove that if the critical point set of R is not contained in the support of a , then I_\hbar is rapidly decreasing in \hbar :

Lemma B.1 *If $dR \neq 0$ on $\text{Supp}(a)$, then $I_{\hbar} = O(\hbar^\infty)$ as $\hbar \rightarrow 0$.*

Proof. Suppose for the moment that $R_{x_1} = \partial R / \partial x_1 \neq 0$ on $\text{Supp}(a)$. Then

$$e^{iR/\hbar} a = -i\hbar \frac{a}{R_{x_1}} \frac{\partial}{\partial x_1} e^{iR/\hbar},$$

and so integration by parts with respect to the x_1 -variable gives

$$|I_{\hbar}| = \hbar \left| \int_{\mathbb{R}^n} e^{iR/\hbar} \frac{\partial}{\partial x_1} \left(\frac{a}{R_{x_1}} \right) |dx| \right|,$$

implying that I_{\hbar} is $O(\hbar)$. Our assertion follows by noting that $(a/Q_{x_1})_{x_1} \in C_0^\infty(\mathbb{R}^n)$ and repeating the same argument.

For the general case, we can use a partition of unity to break up $\text{Supp}(a)$ into finitely many domains as above and then applying the same argument (with x_1 possibly replaced by another coordinate) to each piece. □

The upshot of this lemma is that the main (asymptotic) contribution to the integral I_{\hbar} must come from the critical points of R .

Lemma B.2 *If the quadratic form Q is nondegenerate, then for each nonnegative integer K ,*

$$\int_{\mathbb{R}^n} e^{iQ/\hbar} a |dx| = (2\pi\hbar)^{n/2} \frac{e^{i\pi \text{sgn}(Q)/4}}{|\det_{|dx|} T|^{1/2}} \sum_{k=0}^K \frac{1}{k!} (D^k a)(0) \hbar^k + O(\hbar^{K+1+n/2}),$$

where $T : \mathbb{R}^n \rightarrow (\mathbb{R}^n)^*$ is the self-adjoint map associated to Q and D is the second-order differential operator given by

$$D = \frac{i}{2} \sum_{j,k} T_{jk}^{-1} \frac{\partial}{\partial x_j} \frac{\partial}{\partial x_k}.$$

Proof. From [32, Vol.1, Thm.7.6.1], we recall that for $\xi \in (\mathbb{R}^n)^*$,

$$\int_{\mathbb{R}^n} e^{-i\langle x, \xi \rangle} e^{iQ(x)/\hbar} dx = (2\pi\hbar)^{n/2} \frac{e^{i\pi \text{sgn}(Q)/4} e^{-i\hbar Q^*(\xi)}}{|\det_{|dx|} T|^{1/2}},$$

where $Q^*(\xi) = \langle T^{-1}\xi, \xi \rangle / 2$ and the determinant $\det_{|dx|} T$ is defined as in Appendix A. Consequently, the asymptotic Fourier transform of the function $x \mapsto e^{iQ(x)/\hbar}$ equals

$$\xi \mapsto \frac{e^{i\pi \text{sgn}(Q)/4} e^{-Q^*(\xi)/\hbar}}{|\det_{|dx|} T|^{1/2}}.$$

Combining this expression with the asymptotic Parseval formula, we obtain

$$I_{\hbar} = \int_{\mathbb{R}^n} e^{iQ/\hbar} a |dx| = \frac{e^{i\pi \cdot \text{sgn}(Q)/4}}{|\det_{|dx|} T|^{1/2}} \int_{(\mathbb{R}^n)^*} e^{-iQ^*/\hbar} \overline{\mathcal{F}_{\hbar} a}(\xi) d\xi.$$

A simple computation shows that $\overline{\mathcal{F}_\hbar a}(\xi) = \mathcal{F}_\hbar a(-\xi)$; the change of variables $\xi \mapsto -\xi$ doesn't affect the integral, and consequently we have

$$I_\hbar = \frac{e^{i\pi \cdot \text{sgn}(Q)/4}}{|\det_{|dx}| T|^{-1/2}} \int_{(\mathbb{R}^n)^*} e^{-iQ^*(\xi)/\hbar} \mathcal{F}_\hbar a(\xi) d\xi.$$

To deal with the integral on the right, we use the the Taylor series expansion (with remainder) of $e^{-iQ^*/\hbar}$ to write

$$\begin{aligned} \int_{(\mathbb{R}^n)^*} e^{-iQ^*/\hbar} \mathcal{F}_\hbar a(\xi) d\xi &= \sum_{k=0}^{\infty} \frac{1}{k!} \left(\frac{-i}{\hbar} \right)^k \int_{(\mathbb{R}^n)^*} (Q^*(\xi))^k \mathcal{F}_\hbar a(\xi) d\xi \\ &= (2\pi\hbar)^{n/2} \sum_{k=0}^K \frac{1}{k!} \left(\frac{-i}{\hbar} \right)^k \mathcal{F}_\hbar^{-1}((Q^*)^k \mathcal{F}_\hbar a)(0) + O(\hbar^{K+1+\frac{n}{2}}) \\ &= (2\pi\hbar)^{n/2} \sum_{k=0}^K \frac{1}{k!} (D^k a)(0) \hbar^k + O(\hbar^{K+1+\frac{n}{2}}), \end{aligned}$$

the last expression on the right following from the asymptotic Fourier inversion formula. □

We now wish to apply this lemma to evaluate integrals of the form

$$I_\hbar = \int_M e^{iR/\hbar} \sigma,$$

where M is a smooth n -manifold equipped with a compactly supported density σ , and $R: M \rightarrow \mathbb{R}$ is a smooth function. To this end, we require the following two lemmas. Recall that the hessian of $R: M \rightarrow \mathbb{R}$ at a critical point $p \in M$ is a well-defined self-adjoint linear map $R''(p): T_p M \rightarrow T_p^* M$. The critical point p is called nondegenerate if $R''(p)$ is an isomorphism. In this case, the function R has the following normal form near p .

Morse Lemma . *If p is a nondegenerate critical point of a function $R: M \rightarrow \mathbb{R}$, then there exists a nondegenerate quadratic form Q on \mathbb{R}^n and an embedding $g: U \rightarrow M$, where U is a neighborhood of 0 in \mathbb{R}^n , such that $g(0) = p$ and*

$$(R \circ g)(x) = R(p) + Q(x).$$

□

(This theorem is a special case of the Parametrized Morse Lemma, proven in Section 4.3). If $g: M' \rightarrow M$ is any diffeomorphism, then the “change of variables” formula states that

$$\int_{M'} g^* (e^{iR/\hbar} \sigma) = \int_M e^{iR/\hbar} \sigma.$$

The same role in the stationary phase formula will be played by the following lemma.

Lemma B.3 *In the notation above, suppose that $p \in M$ is a critical point of R and set $p' = g^{-1}(p)$. If σ' is any density on M' such that $\sigma'_{p'} = (g^*\sigma)_p$, then*

$$|\det_{\sigma'}(R \circ g)''(p')| = |\det_{\sigma} R''(p)|.$$

□

Finally, we mention that definition of $|\det_{\sigma}(R)''(p)|$ implies furthermore that

$$f(p) \cdot |\det_{f\sigma} R''(p)|^{1/2} = |\det_{\sigma} R''(p)|^{1/2}$$

for any function f on M . Combining these observations, we obtain

Principle of Stationary Phase . *Let M be a smooth n -manifold and $\sigma \in |\Omega|_c M$. If $R: M \rightarrow \mathbb{R}$ has only nondegenerate critical points p_j , $j = 1, \dots, k$ in $\text{Supp}(\sigma)$, then*

$$\int_M e^{iR/\hbar} \sigma = (2\pi\hbar)^{n/2} \sum_{j=1}^k \frac{e^{iR(p_j)/\hbar} e^{i\pi \cdot \text{sgn}(R''(p_j))/4}}{|\det_{\sigma}(R''(p_j))|^{1/2}} + O(\hbar^{1+n/2}).$$

□

C Čech cohomology

This appendix will give some of the basic definitions of Čech cohomology for manifolds and describe its relation to deRham cohomology. A more general treatment is available in [14].

An open cover $\mathfrak{U} = \{U_{\alpha}\}_{\alpha \in I}$ of a manifold M is said to be *good* if every intersection of finitely many members of \mathfrak{U} is either contractible or empty. For ease of notation, we will denote by $U_{\alpha_0 \dots \alpha_k}$ the intersection $\bigcap_{i=0}^k U_{\alpha_i}$.

If Γ is an abelian group, a Γ -valued Čech cochain with respect to the cover \mathfrak{U} is then a rule which assigns an element $c_{\alpha_0, \dots, \alpha_k}$ of Γ to every list $(\alpha_0, \dots, \alpha_k)$ for which the intersection $U_{\alpha_0 \dots \alpha_k}$ is nonempty. The group of all such cochains is denoted $C_{\mathfrak{U}}^k(M, \Gamma)$, and a coboundary operator

$$\delta^k : C_{\mathfrak{U}}^k(M, \Gamma) \rightarrow C_{\mathfrak{U}}^{k+1}(M, \Gamma)$$

is defined by

$$\delta^k(c)(\alpha_0, \dots, \alpha_{k+1}) = \sum_{j=0}^{k+1} (-1)^j c(\alpha_0, \dots, \widehat{\alpha}_j, \dots, \alpha_{k+1}).$$

where the symbol $\widehat{}$ indicates which member of the list is to be deleted. The groups of degree- k Čech cocycles and coboundaries are defined respectively by

$$\check{Z}_{\mathfrak{U}}^k(M; \Gamma) = \ker(\delta^k) \quad \check{B}_{\mathfrak{U}}^k(M; \Gamma) = \text{im}(\delta^{k-1}),$$

and the k -th Čech cohomology group of M with coefficients in Γ and relative to the covering \mathfrak{U} is the quotient

$$\check{H}_{\mathfrak{U}}^k(M; \Gamma) = \check{Z}_{\mathfrak{U}}^k(M; \Gamma) / \check{B}_{\mathfrak{U}}^k(M; \Gamma).$$

Now consider a closed k -form $\omega \in Z_{DR}^k(M)$. Since each U_α is contractible, there exists on each U_α a $(k-1)$ -form φ_α satisfying $d\varphi_\alpha = \omega$. For any indices α, β , we have

$$d(\varphi_\alpha - \varphi_\beta) = 0$$

on $U_{\alpha\beta}$. Since $U_{\alpha\beta}$ is itself contractible, there exist $(k-2)$ -forms $\psi_{\alpha\beta}$ defined on $U_{\alpha\beta}$ such that

$$d\psi_{\alpha\beta} = \varphi_\alpha - \varphi_\beta$$

and

$$d(\psi_{\alpha\beta} + \psi_{\beta\gamma} - \psi_{\alpha\gamma}) = 0$$

on the set $U_{\alpha\beta\gamma}$ for any indices α, β, γ . Continuing in this way, we see that ω determines a Čech k -cocycle with coefficients in \mathbb{R} . This association defines for each $k \in \mathbb{Z}^+$ a homomorphism

$$Z_{DR}^k(M) \xrightarrow{w} \check{H}_{\mathfrak{U}}^k(M; \mathbb{R}).$$

Theorem C.1 ([61]) *The map w induces an isomorphism between the deRham cohomology of M and the Čech cohomology of M with real coefficients.*

One consequence of this theorem is that $\check{H}_{\mathfrak{U}}^k(M; \Gamma)$ does not depend on the choice of \mathfrak{U} .

A group homomorphism $\Gamma \rightarrow \Gamma'$ induces a homomorphism

$$\check{H}_{\mathfrak{U}}^k(M; \Gamma) \rightarrow \check{H}_{\mathfrak{U}}^k(M; \Gamma')$$

in the obvious way. Of particular interest in these notes is the subgroup $\mathbb{Z}_{\hbar} = 2\pi\hbar\mathbb{Z}$ of \mathbb{R} .

D Principal \mathbb{T}_{\hbar} bundles

In this appendix, we record some standard facts about principal bundles over paracompact manifolds, referring to [17] for more details. Throughout this section, denote by \mathbb{Z}_{\hbar} the group $2\pi\hbar \cdot \mathbb{Z}$ and set $\mathbb{T}_{\hbar} = \mathbb{R}/\mathbb{Z}_{\hbar}$.

A **principal \mathbb{T}_{\hbar} bundle** over a manifold P is a locally trivial \mathbb{T}_{\hbar} bundle $Q \xrightarrow{\pi} P$ together with a nonsingular, fiber-preserving action $\mathbb{T}_{\hbar} \times Q \rightarrow Q$. Two principal \mathbb{T}_{\hbar} bundles $Q \xrightarrow{\pi} P$ and $Q' \xrightarrow{\pi'} P$ are said to be isomorphic if there exists a smooth map $f: Q \rightarrow Q'$ which is equivariant with respect to the \mathbb{T}_{\hbar} actions, i.e. $f(a \cdot p) = a \cdot f(p)$ for all $a \in \mathbb{T}_{\hbar}$ and $p \in Q$, and satisfies $\pi = \pi' \circ f$.

Local triviality of a \mathbb{T}_{\hbar} fiber bundle $Q \xrightarrow{\pi} P$ implies that for any good cover \mathfrak{U} of P , there exist homeomorphisms $h_j: U_j \times \mathbb{T}_{\hbar} \rightarrow \pi^{-1}(U_j)$ such that $h_j(x, t+s) = s \cdot h_j(x, t)$ and $\pi(h_j(x, t)) = x$ for all $(x, t) \in U_j \times \mathbb{T}_{\hbar}$ and $s \in \mathbb{T}_{\hbar}$. These maps give rise to the **transition functions** $g_{jk}: U_{jk} \rightarrow \mathbb{T}_{\hbar}$ of Q , defined by the requirement that

$$h_j(x, t) = h_k(x, t + g_{jk}(x))$$

for all $x \in U_{jk}$. This equation implies that for each $x \in U_{ijk}$,

$$h_i(x, t) = h_i(x, t + g_{ij}(x) + g_{jk}(x) + g_{ki}(x)),$$

and so the transition functions satisfy the cocycle condition $g_{ij} + g_{jk} + g_{ki} = 0 \pmod{\mathbb{Z}_\hbar}$. If $\tilde{g}_{jk} : U_{jk} \rightarrow \mathbb{R}$ is any lift of g_{jk} , then the numbers

$$c_{ijk} = \tilde{g}_{ij} + \tilde{g}_{jk} + \tilde{g}_{ki}$$

are therefore elements of \mathbb{Z}_\hbar which define a Čech cocycle $[c_{ijk}]$. The corresponding class $[Q] \in \check{H}^2(P; \mathbb{Z}_\hbar)$ is known as the **Chern class** of Q . The fundamental theorem describing this space is the following (see [17, Cor.2.1.4] for a proof).

Theorem D.1 *Two principal \mathbb{T}_\hbar bundles over P are isomorphic if and only if their Chern classes are equal. Moreover, the assignment $Q \mapsto [Q]$ induces a bijective map from the space of isomorphism classes of principal \mathbb{T}_\hbar bundles over P to $\check{H}^2(P; \mathbb{Z}_\hbar)$.*

Corresponding to the abelian group structure on $\check{H}^2(P; \mathbb{Z}_\hbar)$ are the following operations on principal \mathbb{T}_\hbar bundles. If $Q \rightarrow P$ is a principal \mathbb{T}_\hbar bundle having transition functions $\{g_{jk}\}$ with respect to some good cover of P , then the **inverse** of Q is defined as the principal \mathbb{T}_\hbar bundle $-Q$ over P obtained from the transition functions $\{-g_{jk}\}$. Similarly, if Q, Q' are principal \mathbb{T}_\hbar bundles over P with transition functions $\{g_{jk}\}$ and $\{g'_{jk}\}$ respectively, then the **product** of Q and Q' is defined as the principal \mathbb{T}_\hbar bundle $Q \times^P Q'$ over P having transition functions $\{g_{jk} + g'_{jk}\}$. From the definitions above, it follows easily that the Chern classes of inverses and products of principal \mathbb{T}_\hbar bundles are given by $[-Q] = -[Q]$ and $[Q \times^P Q'] = [Q] + [Q']$. On the level of the bundles themselves, we can describe the product $Q \times^P Q'$ as the quotient of the usual fiber-product $Q \times_P Q'$ (which in this case is a $\mathbb{T}_\hbar \times \mathbb{T}_\hbar$ -bundle over the base), modulo the anti-diagonal action of \mathbb{T}_\hbar , i.e. $t \cdot (p, p') = (t \cdot p, -t \cdot p')$.

\mathbb{T}_\hbar bundles with connection

The infinitesimal generator of the \mathbb{T}_\hbar action on a principal bundle Q is a vector field X on Q defined by the equation

$$X(p) = \left. \frac{d}{dt} \right|_{t=0} t \cdot p.$$

A **connection** on Q is a \mathbb{T}_\hbar -invariant 1-form φ on Q such that $\varphi(X) = 1$. In terms of a good cover \mathfrak{U} of P , the form φ satisfies $h_j^* \varphi = d\sigma + \pi^* \varphi_j$, where $d\sigma$ denotes the usual form on \mathbb{T}_\hbar and the φ_j are 1-forms on the U_j satisfying

$$\varphi_j - \varphi_k = d\tilde{g}_{jk},$$

where \tilde{g}_{jk} are again \mathbb{R} -valued lifts of the transition functions of Q . The **curvature** of the connection φ is the unique closed 2-form ω on P such that

$$d\varphi = \pi^* \omega.$$

From the compatibility condition for the φ_j , it follows that the Chern class of Q is the Čech representative of the deRham cohomology class $[\omega]$.

Theorem D.2 *A closed 2-form ω on a manifold P is the curvature form of a connection on a principal \mathbb{T}_\hbar bundle Q over P if and only if $\langle \omega, a \rangle \in \mathbb{Z}_\hbar$ for any $a \in H_2(P; \mathbb{Z})$.*

Proof. Most published proofs of this result (e.g. [37]) use Čech cohomology and the deRham isomorphism. We prefer to give the following direct proof by P.Iglesias; see [33] for further details. Let $A(P, p_0)$ denote the space of smooth paths $\gamma: [0, 1] \rightarrow P$ such that $\gamma(0) = p_0$, and let $e: A(P, p_0) \rightarrow P$ be the endpoint map $e(\gamma) = \gamma(1)$. Since the interval $[0, 1]$ is contractible, there exists a natural contraction of $A(P, p_0)$ onto the constant map $[0, 1] \rightarrow p_0$. If Y_t is the vector field which generates this contraction, we define the 1-form $K\omega$ on $A(P, p_0)$ by

$$K\omega = \int_0^1 (Y_t \lrcorner e^*\omega) dt.$$

Then $d\sigma + K\omega$ is a connection form on the product $A(P, p_0) \times \mathbb{T}_{\hbar}$ having curvature $e^*\omega$.

To complete the proof, we will define the \mathbb{T}_{\hbar} bundle (Q, φ) over P with curvature ω as an appropriate quotient of $A(P, p_0) \times \mathbb{T}_{\hbar}$. For this purpose, we call two elements $\gamma, \gamma' \in A(P, p_0)$ *homologous* if their difference is the boundary of a singular 2-chain σ in P . The quotient of $A(P, p_0)$ by this equivalence relation is the covering space \widehat{P} of P corresponding to the commutator subgroup of $\pi_1(P)$, and thus, $H_1(\widehat{P}, \mathbb{Z}_{\hbar}) = 0$. An equivalence relation on the product $A(P, p_0) \times \mathbb{T}_{\hbar}$ is then defined by the condition that $(\gamma, t) \sim (\gamma', t')$ if γ, γ' are homologous and

$$t - t' = \int_{\sigma} \omega,$$

where $\partial\sigma = \gamma - \gamma'$. The quotient of $A(P, p_0) \times \mathbb{T}_{\hbar}$ by this equivalence relation is a principal \mathbb{T}_{\hbar} bundle \widehat{Q} over \widehat{P} with connection $\widehat{\varphi}$ having curvature $\pi^*\omega$, where $\pi: \widehat{P} \rightarrow P$ is the natural projection.

If $H_1(P; \mathbb{Z}_{\hbar}) = 0$, then $\widehat{P} = P$, and the proof is complete. Otherwise, let s be any map from $H_1(P; \mathbb{Z}_{\hbar})$ into the space of loops based at p_0 which assigns a representative to each homology class. A \mathbb{T}_{\hbar} -valued group cocycle on $H_1(P; \mathbb{Z}_{\hbar})$ is then defined by

$$\phi(h, h') = \int_{\sigma} \omega,$$

where $\partial\sigma = s(h + h') - (s(h) + s(h'))$. Since this cocycle is symmetric, it defines a central extension Γ of $H_1(P; \mathbb{Z}_{\hbar})$ which acts naturally on \widehat{Q} by

$$(h, \tau)[\gamma, z] = [s(h) \cdot \gamma, z + \tau].$$

Since \mathbb{T}_{\hbar} is divisible, the extension Γ is isomorphic to the product $\mathbb{T}_{\hbar} \times H_1(P; \mathbb{Z}_{\hbar})$, and any choice of isomorphism defines an action of $H_1(P; \mathbb{Z}_{\hbar})$ on \widehat{Q} which preserves the connection $\widehat{\varphi}$. The quotient of \widehat{Q} is the desired principal \mathbb{T}_{\hbar} bundle Q over P .

□

A connection φ on a principal \mathbb{T}_{\hbar} bundle $Q \xrightarrow{\pi} P$ induces a connection on the inverse $-Q$ of Q whose local representatives are of the form $d\sigma - \varphi_j$, where φ is locally represented by $d\sigma + \varphi_j$. Similarly, connections φ and φ' on the \mathbb{T}_{\hbar} bundles Q, Q' over P induce a connection $\varphi + \varphi'$ on the product $Q \times^P Q'$ defined locally by $d\sigma + \varphi_j + \varphi'_j$.

Two principal \mathbb{T}_\hbar bundles with connection $(Q, \varphi), (Q', \varphi')$ over P are said to be isomorphic provided that there exists an isomorphism $f: Q \rightarrow Q'$ of the underlying principal bundles which satisfies $f^*\varphi' = \varphi$. A check of the definition shows that $(Q, \varphi), (Q', \varphi')$ are isomorphic if and only if $(Q \times^P -Q', \varphi - \varphi')$ is isomorphic to the trivial bundle $(P \times \mathbb{T}_\hbar, dS)$. Since curvature is obviously invariant under isomorphisms of principal \mathbb{T}_\hbar bundles with connection, the classification of such bundles reduces to the classification of flat connections on trivial \mathbb{T}_\hbar bundles over P .

A connection φ on a principal \mathbb{T}_\hbar bundle $Q \xrightarrow{\pi} P$ is said to be **flat** if $d\varphi = 0$. In this case, the local representatives $h_j^*\varphi = d\sigma + \varphi_j$ have the property that $\varphi_j = d\tilde{s}_j$ for functions $\tilde{s}_j: U_j \rightarrow \mathbb{R}$. Denoting by s_j the composition of \tilde{s}_j with the projection $\mathbb{R} \rightarrow \mathbb{T}_\hbar$, we find that the functions $s_{jk} = s_j - s_k$ define transition functions for a trivial \mathbb{T}_\hbar principal bundle over P . On the other hand, the compatibility condition for the φ_j implies that $s_{jk} - g_{jk}$ is constant for each j, k , and so there exists an isomorphism $f: Q \rightarrow P \times \mathbb{T}_\hbar$. It is easy to check that if φ_0 is the trivial connection on $P \times \mathbb{T}_\hbar$, then

$$\varphi - f^*\varphi_0 = \pi^*\beta$$

for some closed 1-form β on P . The cohomology class $[\beta]$ induced by this form in $\check{H}^1(P; \mathbb{T}_\hbar)$ is called the **holonomy** of the (flat) connection φ .

Theorem D.3 *Two flat connections φ, φ' on the trivial principal \mathbb{T}_\hbar bundle $Q = P \times \mathbb{T}_\hbar$ over P are isomorphic if and only if they have equal holonomy. Moreover, the map $(Q, \varphi) \mapsto [\beta_\varphi]$ induces a bijection from the space of isomorphism classes of flat connections on the trivial \mathbb{T}_\hbar bundle with $\check{H}^1(P; \mathbb{T}_\hbar)$.*

A section s of a principal \mathbb{T}_\hbar bundle Q over P with connection φ is called **parallel** provided that $s^*\varphi = 0$. The map $s_Q: Q \rightarrow Q$ associated to a parallel section defines an isomorphism of (Q, φ) with the trivial \mathbb{T}_\hbar bundle equipped with the trivial connection $\varphi = d\sigma$. Thus, we have:

Corollary D.4 *A principal \mathbb{T}_\hbar bundle Q over P with a flat connection admits a parallel section if and only if (Q, φ) has zero holonomy.*

□

Associated line bundles

A representation $\rho: \mathbb{T}_\hbar \rightarrow U(1)$ enables us to associate to any principal \mathbb{T}_\hbar bundle $Q \rightarrow P$ a complex line bundle $E \rightarrow P$ defined explicitly as the quotient of $Q \times \mathbb{C}$ by the \mathbb{T}_\hbar action

$$t \cdot (p, z) = (t \cdot p, \rho^{-1}(t)z).$$

The space of functions $g: Q \rightarrow \mathbb{C}$ satisfying the condition

$$g(t \cdot p) = \rho^{-1}(t)g(p)$$

is identified with the space of sections of E by the assignment $g \mapsto s_G$, where the section $s_g: P \rightarrow E$ is defined by the requirement that

$$s_g(x) = [(p, g(p))],$$

for any element p of $\pi^{-1}(x)$.

References

- [1] Abraham, R., and Marsden, J.E., *Foundations of Mechanics*, 2nd edition, Benjamin/Cummings, Reading, 1978.
- [2] Arnol'd, V.I., *Mathematical Methods of Classical Mechanics*, Graduate Texts in Math. **60**, Springer-Verlag, New York, 1978.
- [3] Arnol'd, V.I., On a characteristic class entering into conditions of quantization, *Funct. Anal. Appl.* **1** (1967) 1-13.
- [4] Arnol'd, V.I., *Catastrophe theory* (3rd edition), Springer-Verlag, New York, 1992.
- [5] Arnol'd, V.I., *Symplectic geometry and topology*, preprint 1993.
- [6] Arnol'd, V.I., and Givental', A.B., *Symplectic geometry*, *Dynamical systems IV*, V.I. Arnol'd and S.P. Novikov, eds., Encyclopedia of Mathematical Sciences, Vol. 4, Springer-Verlag, New York, 1990, pp. 1-136.
- [7] Atiyah, M.F., *The Geometry and Physics of Knots*, Cambridge Univ. Press, Cambridge, 1990.
- [8] Audin, M., *The Topology of Torus Actions on Symplectic Manifolds*, Progress in Math. vol. 93, Birkhäuser, Basel, 1991.
- [9] Bayen, F., Flato, M., Fronsdal, C., Lichnerowicz, A., and Sternheimer, D., Deformation theory and quantization, I and II, *Ann. Phys.* **111**, (1977), 61-151.
- [10] Berezin, F.A., Some remarks about the associated envelope of a Lie algebra, *Funct. Anal. Appl.* **1**, (1967), 91-102.
- [11] Berezin, F.A., Quantization, *Math USSR Izv.* **8** (1974), 1109-1165.
- [12] Besse, A. *Manifolds All of Whose Geodesics are Closed*, Ergebnisse der Math. 93 Springer-Verlag, Berlin, 1978.
- [13] Bohm, D., *Quantum Theory*, Prentice-Hall, New York, 1951.
- [14] Bott, R. and Tu, L., *Differential Forms in Algebraic Topology*, Springer-Verlag, New York, 1982.
- [15] Boutet de Monvel, L., and Guillemin, V., *The Spectral Theory of Toeplitz Operators*, Annals. of Math. Studies **99**, Princeton University Press, Princeton, 1981.
- [16] Brüning, J. and Guillemin, V., *Mathematics Past and Present: Fourier Integral Operators*, Springer-Verlag, Berlin, 1994.
- [17] Brylinski, J.-L., *Loop Spaces, Characteristic Classes, and Geometric Quantization*, Progress in Mathematics vol. 107, Birkhäuser, Boston, 1992.

- [18] Connes, A., *Noncommutative Geometry*, Academic Press, San Diego, 1995.
- [19] De Wilde, M., and Lecomte, P., *Formal deformations of the Poisson Lie algebra of a symplectic manifold and star-products. Existence, equivalence, derivations*, in M. Hazewinkel and M. Gerstenhaber, eds., *Deformation Theory of Algebras and Structures and Applications*, Kluwer Acad. Pub., Dordrecht (1988), 897-960.
- [20] Dirac, P. A. M., *The Principles of Quantum Mechanics*, Clarendon Press, Oxford, 1930.
- [21] Duistermaat, J. J., *Fourier Integral Operators*, Courant Institute, NYU, New York, 1973. (New edition to be published by Birkhäuser Boston in 1995.)
- [22] Ehresmann, C., Les connexions infinitésimales dans un espace fibré différentiable, *Colloque de Topologie (Espaces Fibrés) tenu a Bruxelles du 5 au 8 juin, 1950*, Georges Thone, Liège and Masson et Cie., Paris, 1951, pp. 29-55.
- [23] Fedosov, B., A simple geometrical construction of deformation quantization, *J. Diff. Geom.* **40** (1994) 213-238.
- [24] Folland, G.B., *Harmonic Analysis in Phase Space*, Annals of mathematics studies **122**, Princeton University Press, 1989.
- [25] Fomenko, A.T., *Symplectic Geometry*, Gordon and Breach, New York, 1988.
- [26] Golubitsky, M., Guillemin, V., *Stable Mappings and Their Singularities*, Graduate Texts in Math., vol. 14, Springer-Verlag, (1973).
- [27] Gompf, R., New constructions of symplectic manifolds, *Ann. Math.*, to appear.
- [28] Guillemin, V., and Sternberg, S., *Geometric Asymptotics*, Math. Surveys **14**, Amer. Math. Soc., Providence, 1977.
- [29] Guillemin, V., and Sternberg, S., *Symplectic Techniques in Physics*, Cambridge University Press, Cambridge, 1984.
- [30] Helffer, B., Semi-classical analysis for the Schrodinger operator and applications, *Lecture Notes in Math.* **1336** (1988).
- [31] Hörmander, L., Fourier integral operators I., *Acta Math.* **127** (1971) 79-183.
- [32] Hörmander, L., *The Analysis of Linear Partial Differential Operators*, Vols. 1-4, Springer-Verlag, New York, 1983-90.
- [33] Iglesias, P., La trilogie du moment, *Ann. Inst. Fourier* (to appear).
- [34] Karasev, M.V. and Maslov, V.P., *Nonlinear Poisson Brackets : Geometry and Quantization* Translations of mathematical monographs vol. 119, American Mathematical Society, Providence, R.I., 1993.

- [35] Kirillov, Geometric quantization, *Dynamical Systems IV*, V.I. Arnol'd and S.P. Novikov, eds., Encyclopedia of Mathematical Sciences, Vol. 4, Springer-Verlag, New York, 1990, pp. 137-172.
- [36] Klingenberg, W., *Riemannian Geometry*, deGruyter Studies in Math., Berlin, 1982.
- [37] Kostant, B., Quantization and unitary representations, *Lecture Notes in Math.* **170** (1970), 87-208.
- [38] Libermann, P., and Marle, C.-M., *Symplectic Geometry and Analytical Mechanics*, Reidel, Dordrecht, 1987.
- [39] Lie, S., *Theorie der Transformationsgruppen*, (Zweiter Abschnitt, unter Mitwirkung von Prof. Dr. Friedrich Engel), Leipzig, Teubner, 1890.
- [40] McDuff, D., Examples of symplectic structures, *Invent. Math.* **89** (1987), 13-36.
- [41] Mackey, G., *Mathematical Foundations of Quantum Mechanics*, W.A.Benjamin, Reading, 1963.
- [42] Marsden, J., *Lectures on Mechanics*, London Math. Soc. Lecture Note Series 174, Cambridge Univ. Press, Cambridge, 1992.
- [43] Maslov, V.P., *Theory of Perturbations and Asymptotic Methods*, Moskov. Gos. Univ., Moscow, 1965.
- [44] Mélotte, D., Invariant deformations of the Poisson Lie algebra of a symplectic manifold and star products, in M. Hazewinkel and M. Gerstenhaber, eds., *Deformation Theory of Algebras and Structures and Applications*, Kluwer Acad. Pub., Dordrecht (1988), 961-972.
- [45] Moser, J., On the volume elements on a manifold, *Trans. Amer. Math. Soc.* **120** (1965), 280-296.
- [46] Nazaikinskii, V., Sternin, B., and Shatalov, V., *Methods of Noncommutative Analysis, Theory and Applications*, Berlin, New York, W. de Gruyter (to appear).
- [47] Olver, F.W.J., *Asymptotics and Special Functions*, Academic Press, New York, 1974.
- [48] Omohundro, S., *Geometric Perturbation Theory in Physics*, World Scientific, Singapore, 1986.
- [49] Omori, I., *Infinite-dimensional Lie Transformation Groups*, Lecture Notes in Mathematics, vol.427, Springer-Verlag, New York, 1975.
- [50] Rieffel, M., Deformation quantization of Heisenberg manifolds, *Comm. Math. Phys.* **122** (1989), 531-562.

- [51] Robert, D., *Autour de l'Approximation Semi-classique*, Progress in Mathematics, vol. 68, Birkhäuser, Boston, 1987.
- [52] Segal, I., Quantization of nonlinear systems, *J. Math. Phys.* **1** (1960), 468-488.
- [53] Sniatycki, J., *Geometric Quantization and Quantum Mechanics*, Springer-Verlag, New York, 1980.
- [54] Souriau, J.M., *Structure des Systèmes Dynamiques*, Dunod, Paris, 1970.
- [55] Souriau, J.M., Groupes différentiels, *Lecture Notes in Math.* **836** (1980), 91-128.
- [56] Treves, F., *Introduction to Pseudodifferential and Fourier Integral Operators..* Vols I,II. Plenum Press, New York, 1980.
- [57] Tulczyjew, W.M., The Legendre transform, *Ann. Inst. H. Poincaré*, **27** (1977) 101-114.
- [58] Unterberger, A., Quantification et analyse pseudo-différentielle, *Ann. Sci. Éc. Norm. Sup.* **21** (1988), 133-158.
- [59] Vaisman, I., *Symplectic Geometry and Secondary Characteristic Classes*, Birkhäuser, Boston, 1987.
- [60] Wallach, N. R., *Symplectic Geometry and Fourier Analysis*, Math. Sci. Press, Brookline, 1977.
- [61] Weil, A., Sur les théorèmes de de Rham, *Comm. Math. Helv.* **26** (1952) 119-145.
- [62] Weinstein, A., Symplectic manifolds and their lagrangian submanifolds, *Advances in Math.* **6** (1971), 329-346.
- [63] Weinstein, A., *Lectures on Symplectic Manifolds*, Regional conference series in mathematics **29**, Amer. Math. Soc., Providence, 1977.
- [64] Weinstein, A., Symplectic geometry, *Bulletin A.M.S.* (new series) **5** (1981), 1-13.
- [65] Weinstein, A., Coisotropic calculus and Poisson groupoids, *J. Math. Soc. Japan* **40** (1988), 705-727.
- [66] Weinstein, A., Noncommutative geometry and geometric quantization, in *Symplectic Geometry and Mathematical Physics: actes du colloque in l'honneur de Jean-Marie Souriau*, P.Donato et al. eds., Birkhäuser, Basel, 1991.
- [67] Weinstein, A., Symplectic groupoids, geometric quantization, and irrational rotation algebras, in *Symplectic Geometry, Groupoids, and Integrable Systems*, P.Dazord and A.Weinstein, eds., MSRI Publications, Springer, New York, 1991.
- [68] Weinstein, A., Deformation quantization, *Séminaire Bourbaki*, no. 789, Juin 1994.

- [69] Weinstein, A., and Xu, P., Extensions of symplectic groupoids and quantization, *J. Reine Angew. Math.* **417** (1991), 159-189
- [70] Weyl, H., *The Theory of Groups and Quantum Mechanics*, Dover Publications, New York, 1950.
- [71] Woodhouse, N.M.J., *Geometric Quantization*, 2nd Ed., Oxford U. Press, Oxford, 1992.
- [72] Yano, K., and Ishihara, S., *Tangent and Cotangent Bundles; Differential Geometry* Dekker, New York, 1973.