PQM Supplementary Notes: Spin, topology, SU(2)→SO(3) etc

(away from the syllabus, but interesting material, based on notes by Dr J.M. Evans)

1 Rotations and Non-contractible Loops

Rotations in $\mathbf{R}^3$ can be thought of as the group of real **special orthogonal** matrices

$$\mathrm{SO}(3) = \{\, R \ 3{\times}3 \text{ real matrix } : \ R^T = R^{-1}\,,\ \det R = 1\} \ .$$

Each rotation can also be specified geometrically by an axis $\mathbf{n}$ (a unit vector) and an angle θ which can be combined $\xi = \theta\mathbf{n}$. If we were to choose $0 \le \theta \le \pi$, we would have a parametrization which is almost one-to-one—it is just slightly redundant in that rotations with axes $\pm\mathbf{n}$ and $\theta = \pi$ are the same. Thus, we can picture the space of all rotations as a solid, three-dimensional ball $|\xi| \le \pi$, but with opposite or **antipodal** points on the boundary identified. As θ becomes larger than π, it appears on the other side of the ball. Thus the radius of position in the ball gives the rotation angle, minus some integer factor of π.

This description of the rotation group reveals an important topological property: there are **loops**, or smooth closed paths, which cannot be smoothly shrunk to a point. Consider paths which start and end at the identity rotation, i.e. the centre of the sphere. Figure 1a shows a loop which can obviously be shrunk to a point—it is **contractible**. Figure 1b shows a loop for which the angle of rotation starts at zero, smoothly increases to π (at the point A where it reaches the boundary and reappears at the antipodal point A'), and then continues to increase to 2π, back at the identity. It is clear intuitively that the loop in figure 1b cannot be shrunk to a point while keeping its ends fixed—it is **non-contractible**.

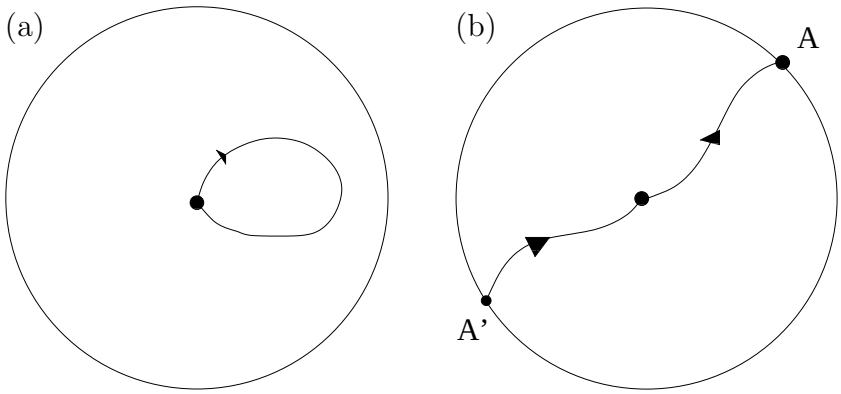

Figure 1: Examples of (a) a contractible loop, (b) a non-contractible loop. The ends of the loop at the origin are fixed.

Now consider the loop in figure 2a, along which the angle of rotation increases from 0 to 4π, requiring the boundary to be crossed twice, once at A, reappearing at A', then again

at B, reappearing at B'. By moving the second antipodal pair B and B', as shown in figures 2b,c, the section of the path between A' and B can be pulled across the boundary, smoothly deforming it to the situation in figure 1a. Thus after **two** complete rotations the situation becomes simple once more. In general, the angle along any closed path must increase by an integer multiple of 2π, and the loop will be contractible if this integer is even but non-contractible if it is odd.

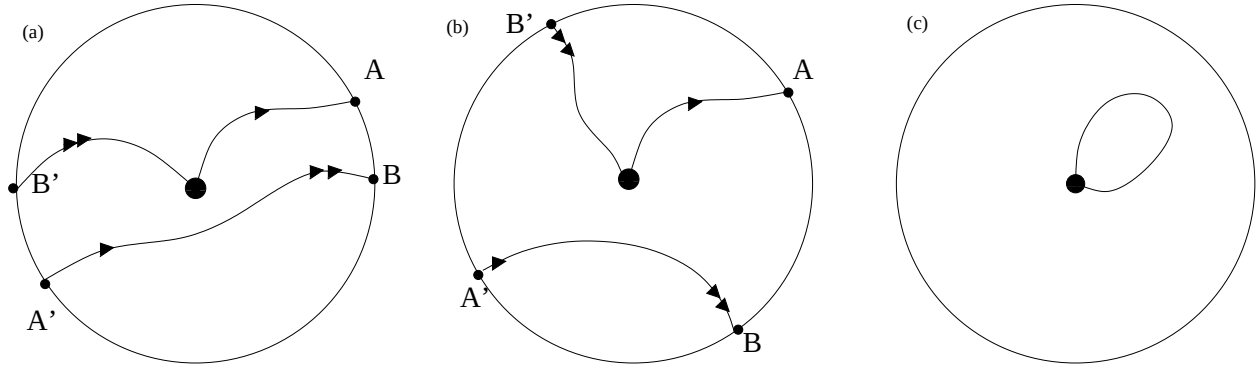

Figure 2: A contractible loop, as shown by the continuous deformations of (a) to (c). The ends of the loop at the origin are fixed.

Why should any of this be important for Quantum Mechanics? We are used to thinking that a rotation by 2π is precisely the same as no rotation at all, but we have seen above that this is not quite right if the rotations arise as part of a smooth family. The structure of Quantum Mechanics is sensitive to this and is capable of keeping track of whether the number of complete rotations we perform is even or odd. To understand how, we must first consider symmetries in a little more generality.

2 Symmetries and Projective Representations

Symmetries in QM are represented by unitary operators: given a group G there is a unitary $U(g)$ for each $g \in G$. There is a freedom or ambiguity in how these operators are defined, however, because states in QM are physically equivalent if they are multiples of one another, so a change

$$U(g) \to \alpha(g)U(g) \tag{1}$$

where $\alpha(g)$ is any complex number of unit modulus, does not affect the physical interpretation of the operator. This freedom is often sufficient to allow us to choose operators which obey

$$U(g_1)U(g_2) = U(g_1 g_2) \qquad \text{for all} \qquad g_1,\, g_2 \in G$$

which is a **representation** of G on the states. More generally, though,

$$U(g_1)U(g_2) = \omega(g_1, g_2)U(g_1 g_2) \qquad \text{for all} \qquad g_1,\, g_2 \in G$$

where the complex number $\omega(g_1, g_2)$ is called a **cocycle**; this is a **projective representation** on states.

For G a continuous group—actually **smooth** in an appropriate sense, a **Lie group**—we would also like $U(g)$ to depend smoothly on g. The topological properties of paths in G then become important when we consider the behaviour of $U(g)$ as the group element g varies around a loop ℓ beginning and ending at some fixed g_0. It is quite possible for $U(g)$ to change smoothly with g along ℓ in such a way that the operators at each end of the loop do not coincide: consistent with the ambiguity (1), they may differ by a change

$$U(g_0) \to \alpha_\ell U(g_0) \tag{2}$$

which depends ℓ. It would obviously be simplest if we could arrange to have $\alpha_\ell = 1$ for all loops, and to investigate this we specialise from here on to the case in which the operators act on a finite-dimensional space of states.

Let N be the dimension of the space of states, so that each $U(g)$ can be regarded as an $N \times N$ matrix. We can use up some of the freedom in (1) to impose the condition that each $U(g)$ has unit determinant. This does not fix things completely, but the residual freedom in (1) or (2) is discrete:

$$\det U(g) = 1 \qquad \Rightarrow \qquad \alpha^N = 1 \tag{3}$$

Notice too that once the determinant condition is imposed, any cocycle must also obey $\omega^N = 1$. We will assume (3) holds henceforth.

If G is **simply-connected**, meaning that **any** closed loop is contractible, then the determinant condition implies that $\alpha_\ell = 1$ for all ℓ. This follows straightforwardly using continuity: α_ℓ is restricted to a discrete set of values by (3) and it must therefore remain unchanged under any smooth variation in ℓ. If, in addition, the loop ℓ can be shrunk to the single point g_0, then we must have the same matrix at the beginning and end of the loop, so $\alpha_\ell = 1$. If G is not simply connected, the same argument still shows that $\alpha_\ell = 1$ for any contractible ℓ, but there may also be non-contractible loops with $\alpha_\ell \neq 1$. In this case, we can at least deduce, in a similar fashion, that $\alpha_\ell = \alpha_{\ell'}$ whenever the loops ℓ and ℓ' are in the same **topological class**, meaning that they can be smoothly deformed into one another. Furthermore, we observe that if loops ℓ and ℓ' are traversed successively, then the change in $U(g_0)$ is $\alpha_\ell \alpha_{\ell'}$, so there are self-consistency constraints.

For a simply-connected group G, we can now give a unique definition of $U(g)$ for all g, by using (3). Setting $U(1_G) = 1$, define $U(g_0)$ by choosing any path from the identity 1_G to g_0 and demanding that $U(g)$ changes smoothly along it. The values along the path are unique (by the determinant condition and continuity) but the end result $U(g_0)$ is also independent of the path, because any two paths can be combined to form a closed loop at g_0 (traversing one path backwards and then the other forwards) which is contractible. With this definition, continuity also ensures that all cocycles are equal to one, so that we have a genuine representation of G.

Carrying out the same construction when G is not simply connected, we will encounter paths from 1 to g_0 which cannot be smoothly deformed into one another and, starting from $U(1_G) = 1$, we will in general obtain different values for $U(g_0)$ depending on the path we take. In this way we are forced to consider **multi-valued** functions on G, just as when defining a continuous square-root in the complex plane. The ambiguity or multi-valuedness in $U(g_0)$ can be resolved only by keeping track of the path used to reach g_0, just as for the complex square root we must keep track of how many times we have encircled the origin to be sure which branch we are on. Such a multi-valued definition inevitably means that non-trivial cocycles appear.

We are at last in a position to understand rotations properly. The group $G = \mathrm{SO}(3)$ is not simply connected, but there are just two topological classes of loops, depending on whether the net angle of rotation is an even or odd multiple of 2π. Any loop ℓ in the first class is contractible, and so has $\alpha_\ell = 1$. Any loop in the second class is non-contractible, but if we traverse ℓ twice the result becomes contractible again, and so $\alpha_\ell^2 = 1$, or $\alpha_\ell = \pm 1$.

The finite-dimensional spaces on which the rotation operators act are nothing but the multiplets of angular momentum j with basis states

$$\{\, |j\, m\rangle : m = j, j{-}1, \ldots, -j \,\}\ , \quad \text{dimension}\ \ N = 2j{+}1\ , \qquad \text{for } j = 0, \tfrac{1}{2}, 1, \tfrac{3}{2}, \ldots\ .$$

From (3) we know that $\alpha_\ell^N = 1$ but from above $\alpha_\ell = \pm 1$. If j is an integer then N is odd and it follows that $\alpha_\ell = 1$ for all loops, whether they are contractible or not. In this case we have a genuine representation of SO(3) on the states. If j is a half-integer then N is even and $\alpha_\ell = -1$ is possible for non-contractible loops. This is exactly the behaviour found by the explicit calculation in lectures for $j = \frac{1}{2}$ states: under a continuous family of rotations through a net angle which is an odd multiple of 2π the states change sign. The discussion above shows that the origin of this unexpected sign is the non-trivial topology of the rotation group—in our new terminology, we have a projective representation of SO(3).

3 The Double Cover SU(2) → SO(3)

It is not too difficult to live with the multi-valuedness of $U(g)$ for half integral spin representations of SO(3), since the ambiguity is just a sign ± 1 in either the operators or the cocycles. Nevertheless it is natural to ask if we can do better.

Returning to symmetries in general, one way of getting around the multi-valuedness of projective representations is to view this as a problem not with the representation but rather with the underlying group. From this standpoint, the true symmetry at work is not exactly G but something closely related to it, the **covering group** $\tilde{G}$. Technically, we require a homomorphism $\tilde{G} \to G$ which is onto and many-to-one, but in just such a way that the multi-valued or projective representations of G descend from genuine, single-valued representations of $\tilde{G}$. Pursuing the analogy with the complex square root, this is like saying that the problem is not with the form of the function, as such, but rather with the domain of its definition. The square root becomes single-valued on a new domain obtained by gluing together two copies of the cut plane, and there is a two-to-one map from the new set to old.

For the rotation group $G = \mathrm{SO}(3)$, the simply-connected covering group is $\tilde{G} = \mathrm{SU}(2)$, a **special unitary** group defined by

$$\mathrm{SU}(2) = \{\ U\ \ 2{\times}2 \text{ complex matrix } :\ U^\dagger = U^{-1}\,,\ \det U = 1\}\ .$$

There is a **double cover**, or two-to-one, surjective homomorphism

$$\Phi :\quad \mathrm{SU}(2) \to \mathrm{SO}(3)\ , \qquad \ker\Phi = \mathbf{Z}_2 = \{\pm 1\}$$

which we will now construct with the help of the Pauli matrices.

Let $\mathbf{H}$ denote the set of traceless, hermitian, $2{\times}2$ complex matrices, which is a real vector space with the standard definitions of addition and scalar multiplication of matrices. Given any $X \in \mathbf{H}$ and any $U \in \mathrm{SU}(2)$ we note that $UXU^\dagger$ is also hermitian and $\mathrm{tr}(UXU^\dagger) =$

$\mathrm{tr}(X) = 0$, by cyclicity of the trace, so each U defines a map $\mathbf{H} \to \mathbf{H}$. This map is clearly linear, and has the property

$$X \mapsto UXU^\dagger \ , \quad Y \mapsto UYU^\dagger \ , \qquad \Rightarrow \qquad \mathrm{tr}(XY) \ \mapsto \ \mathrm{tr}(UXU^\dagger UYU^\dagger) \ = \ \mathrm{tr}(XY) \ . \tag{4}$$

Now recall that the Pauli matrices $\sigma = (\sigma_1, \sigma_2, \sigma_3)$ each belong to $\mathbf{H}$ and that they obey

$$\sigma_i\sigma_j = \delta_{ij} + i\epsilon_{ijk}\sigma_k \ , \qquad \mathrm{tr}(\sigma_i\sigma_j) = 2\delta_{ij} \ . \tag{5}$$

The Pauli matrices provide a convenient basis which we can use to identify $\mathbf{H} \leftrightarrow \mathbf{R}^3$ by associating to any vector $\mathbf{x} \in \mathbf{R}^3$ a matrix $X = \mathbf{x}\cdot\sigma \in \mathbf{H}$, or vice versa. This one-to-one correspondence is an isomorphism of vector spaces, and from (5),

$$\mathbf{x} \leftrightarrow X = \mathbf{x}\cdot\sigma \ , \quad \mathbf{y} \leftrightarrow Y = \mathbf{y}\cdot\sigma \qquad \Rightarrow \qquad \mathbf{x}\cdot\mathbf{y} \ = \ \tfrac{1}{2}\mathrm{tr}(\mathbf{x}\cdot\sigma\,\mathbf{y}\cdot\sigma) \ = \ \tfrac{1}{2}\mathrm{tr}(XY) \ , \tag{6}$$

so there is also a natural way to carry over the standard inner-product from $\mathbf{R}^3$ to $\mathbf{H}$. With the identifications we have just made, a linear map $\mathbf{H} \to \mathbf{H}$ defined by some $U \in \mathrm{SU}(2)$ corresponds to a linear map $\mathbf{R}^3 \to \mathbf{R}^3$ given by some real 3×3 matrix R:

$$\mathbf{x} \mapsto R\mathbf{x} \qquad \text{where} \qquad \mathbf{x}\cdot\sigma \ \mapsto \ U\mathbf{x}\cdot\sigma U^\dagger \ = \ (R\mathbf{x})\cdot\sigma \qquad \text{for all} \ \ \mathbf{x} \ . \tag{7}$$

But by comparing (4) and (6) we see that $(R\mathbf{x})\cdot(R\mathbf{y}) = \mathbf{x}\cdot\mathbf{y}$, so R is orthogonal. Furthermore, $\mathrm{Tr}\{\,(R\mathbf{e}_1)\cdot\sigma\,(R\mathbf{e}_2)\cdot\sigma\,(R\mathbf{e}_3)\cdot\sigma\,\} \ = \ \mathrm{Tr}(\,U\sigma_1U^\dagger U\sigma_2U^\dagger U\sigma_3U^\dagger\,) \ = \ \mathrm{Tr}(\sigma_1\sigma_2\sigma_3)$ implies, using (5), that $\epsilon_{ijk}R_{i1}R_{j2}R_{k3} = \det R = 1$, so $R \in \mathrm{SO}(3)$. The double cover is the map

$$\Phi : \ \ \mathrm{SU}(2) \to \mathrm{SO}(3) \ , \qquad \Phi(U) = R \tag{8}$$

with U and R related by (7). It is easily checked to be a group homomorphism with the kernel $\{\pm1\}$, as claimed above.

Finally, we return to spin half in Quantum Mechanics. The motivation for considering the covering group was to eliminate the need for multi-valued or projective representations. We know (from lectures and example sheets) that the unitary rotation operator for spin half states is easily calculated:

$$U(\theta\mathbf{n}) = \exp(\,-i\tfrac{1}{2}\theta\mathbf{n}\cdot\sigma\,) = \cos\tfrac{1}{2}\theta - i\mathbf{n}\cdot\sigma\sin\tfrac{1}{2}\theta \ . \tag{9}$$

For a rotation R with axis $\mathbf{n}$ and angle θ, this formula gives, explicitly, an element U of $\mathrm{SU}(2)$ with $\Phi(U) = R$. If we change θ by 2π then R is unchanged and U changes sign. But this is of course consistent with $\Phi(-U) = R$, the map Φ being two-to-one. Thus, the example which first led us to observe the multi-valued nature of the rotation operators on half-integral spin states is nothing but the simplest (or defining) representation for SU(2), the covering group of SO(3).

Scattering and Cross-Sections from the Golden Rule

(an alternative derivation of some results in Part II AQM)

A scattering process involves a particle, or beam of particles, being sent towards a target, interacting with it, and then emerging along some new direction. The aim is to calculate the probability for scattering in a given range of final directions, and to understand how

this depends on the energy of the incident particle or beam. Applications are wide-ranging: from α-particles incident on atoms, as in the original Rutherford scattering experiment, to X-rays or neutrons diffracting from crystals, scattering processes have provided crucial evidence concerning the nature of the targets, e.g. the existence of the atomic nucleus, or the lattice structure of crystalline solids.

We can, in principle, consider an initial wave packet, localised in space, which is fired in towards a target and which scatters off it by evolving according to the time-dependent Schrödinger equation. In general, this approach is far too complicated to be tractable, however, and we deal instead with an idealisation where the initial and final states, $|a\rangle = |\psi_{\mathbf{k}}\rangle$ and $|b\rangle = |\psi_{\mathbf{k}'}\rangle$, are eigenstates of the free particle Hamiltonian H_0, with definite momenta $\hbar\mathbf{k}$ and $\hbar\mathbf{k}'$ respectively. The full Hamiltonian

$$H_0 + V \qquad \text{with} \qquad H_0 = \frac{1}{2m}\mathbf{p}^2$$

incorporates the interaction with the target through the potential V which we assume is independent of time and falls off rapidly $V(\mathbf{x}) \to 0$ as $|\mathbf{x}| \to \infty$. Treating the potential by first-order time-dependent perturbation theory, the transition rate from a to b is given by **Fermi's Golden Rule**:

$$\frac{2\pi}{\hbar}\rho(E)\,|\,\langle b|V|a\rangle\,|^2$$

where $\rho(E)$ is the density of final states (with respect to energy) with certain specified parameters, and $E = E_a = E_b$, or $|\mathbf{k}| = |\mathbf{k}'|$ in our case.

Let us choose coordinates with the z-axis along the direction of the incident particle or beam, so $\mathbf{k} = k\hat{\mathbf{z}}$. Outgoing particle states have $\mathbf{k}' = k\mathbf{n}$ for some unit vector $\mathbf{n}$, and we will consider scattering into a small solid angle $\delta\Omega$, corresponding to a small variation in the direction of $\mathbf{n}$. The number of such final states with energy in the range E to $E+\delta E$ is, by definition, $\rho(E,\delta\Omega)\delta E$. We expect this to be proportional to $\delta\Omega$ and this will be confirmed shortly. In addition, we expect the transition rate to our specified set of final states to be proportional to the incident particle flux, as calculated from the initial state.

In the Golden Rule, the initial and final states must be normalised, and we also need to be able to count states to determine $\rho(E,\delta\Omega)$. It is very convenient, therefore, to put the entire system in a large cubical box with sides of length L, parallel to the coordinate axes, so that the momenta and eigenvalues of H_0 become discrete (we could think of this box as the laboratory). Rather than demanding that the wavefunctions vanish on the faces of the box, however, it is simpler to impose **periodic boundary conditions**, which imply that the allowed eigenstates of H_0 and their eigenvalues are

$$\psi_\kappa(\mathbf{x}) = \frac{1}{L^{3/2}}e^{i\kappa\cdot\mathbf{x}} \quad \text{with} \quad E_\kappa = \frac{\hbar^2\kappa^2}{2m} \qquad \text{where} \qquad \kappa = \frac{2\pi}{L}(N_1,N_2,N_3) \quad \text{for} \quad N_i \in \mathbf{Z}\,.$$

It is easy to check

$$\int_{\text{box}} \psi_{\kappa'}(\mathbf{x})^*\psi_\kappa(\mathbf{x})\,d^3\mathbf{x} = \begin{cases} 1 & \text{if } \kappa = \kappa' \\ 0 & \text{else} \end{cases}$$

Our final results must be independent of the choice of boundary conditions and of the size of the box, which we will eventually remove by taking $L \to \infty$.

With the values for κ permitted above, the number of states lying within a certain volume in κ-space is approximately $(L/2\pi)^3$ times the volume, and this approximation becomes

better and better the larger L becomes. We are interested in final states with $\kappa = \mathbf{k}' = k\mathbf{n}$, the direction $\mathbf{n}$ constrained within $\delta\Omega$, and the size of the wave vector in the range k to $k+\delta k$, corresponding to a range of energies E to $E+\delta E$. The volume in κ-space is therefore $k^2\delta k\delta\Omega$ and so the number of states is

$$\Big(\frac{L}{2\pi}\Big)^3 k^2\delta k\,\delta\Omega \;=\; \Big(\frac{L}{2\pi}\Big)^3 \frac{mk}{\hbar^2}\,\delta E\,\delta\Omega \;=\; \rho(E,\delta\Omega)\delta E\ ,$$

which determines the density $\rho(E,\delta\Omega)$.

The Golden Rule also requires us to calculate

$$\langle\psi_{\mathbf{k}'}|V|\psi_{\mathbf{k}}\rangle \;=\; \int \psi_{\mathbf{k}'}(\mathbf{x})^*\,V(\mathbf{x})\,\psi_{\mathbf{k}}(\mathbf{x})\,d^3\mathbf{x} \;=\; \frac{1}{L^3}\int V(\mathbf{x})\,e^{-i\mathbf{q}\cdot\mathbf{x}}\,d^3\mathbf{x} \;=\; \frac{1}{L^3}\,\tilde{V}(\mathbf{q})\ ,$$

where $\hbar\mathbf{q} = \hbar\mathbf{k}' - \hbar\mathbf{k}$ is the **momentum transfer**. Plugging this matrix element and the density into the formula, we get the rate for scattering into solid angle $\delta\Omega$ at energy E:

$$\frac{2\pi}{\hbar}\,\rho(E,\delta\Omega)\,|\langle\psi_{\mathbf{k}'}|V|\psi_{\mathbf{k}}\rangle|^2 \;=\; \frac{1}{(2\pi)^2}\frac{mk}{\hbar^3L^3}\,|\tilde{V}(\mathbf{q})|^2\delta\Omega\ .$$

The last step is to divide this rate by the particle flux $j = \hbar k/mL^3$ corresponding to the initial wavefunction $\psi_{\mathbf{k}}(\mathbf{x})$. The result is the **differential cross-section** $\delta\sigma$ for scattering into the solid angle $\delta\Omega$ at energy E and is usually written, in the limit $\delta\Omega \to 0$,

$$\frac{d\sigma}{d\Omega} = \Big(\frac{m}{2\pi\hbar^2}\Big)^2|\tilde{V}(\mathbf{q})|^2\ , \qquad \mathbf{q} = \mathbf{k}' - \mathbf{k}\ . \tag{$*$}$$

With the rate correctly normalised relative to the incident flux, the resulting cross-section is independent of L and we can remove the box by taking $L \to \infty$.

The formula $(*)$, giving the cross-section to leading order in the potential V, is known as the **Born approximation** for scattering. In applying the Golden Rule we have, of course, used first-order time-dependent perturbation theory. There are other ways of getting the same result—for instance, by modelling the scattering process as a stationary state of the entire Hamiltonian with an appropriate asymptotic form for the wavefunction as $|\mathbf{x}| \to \infty$ (see Part II Applications of Quantum Mechanics). One nice feature of the derivation given here is that it is fairly easily adapted to cases in which the **scatterer**, i.e. the source of the potential V, also undergoes a change of state, e.g. a particle scattering off an atom and causing a transition.

An important application of $(*)$ is to scattering in the screened Coulomb potential

$$V(\mathbf{x}) = \frac{\lambda}{4\pi}\frac{1}{|\mathbf{x}|}e^{-\mu|\mathbf{x}|} \qquad \Rightarrow \qquad \tilde{V}(\mathbf{q}) = \frac{\lambda}{\mathbf{q}^2+\mu^2}\ .$$

The exponential factor is convenient mathematically since it allows a sound calculation of the Fourier transform (it ensures that V tends to zero sufficiently rapidly at large distances). It can also be motivated physically, e.g. if we are considering scattering off an atomic nucleus, it gives a way of representing the screening effects of the surrounding electrons (which should indeed become negligible for $|\mathbf{x}|$ very small, where the nucleus dominates). If θ and ϕ are the usual polar angles (with our choice of z-axis along the direction of the incident beam) then a simple calculation reveals $(\hbar\mathbf{q})^2 = 8mE\sin^2\frac{1}{2}\theta$, and the scattering is independent of

ϕ. This is to be expected for a spherically symmetric potential, since the scattering problem is then axis symmetric about the direction of the incident beam.

To study scattering in a pure Coulomb potential we take $\mu = 0$ to obtain the **Rutherford cross-section**

$$\frac{d\sigma}{d\Omega} = \frac{\lambda^2}{(16\pi E \sin^2 \frac{1}{2}\theta)^2} \ .$$

This is the theoretical basis for understanding the scattering experiments which first proved the existence of a compact, massive nucleus within the atom (α-particles, or Helium nuclei, incident on a thin metal film). There is an historical twist here, however, because the result above is independent of $\hbar$, reflecting the fact that this cross-section can be obtained by a purely **classical** scattering calculation for the Coulomb potential (this is essentially because when $\mu = 0$ the answer is severely restricted by dimensional analysis). So Rutherford was able to get the right answer using Classical Mechanics—but then the concept of the nuclear atom which he introduced led to the development of Quantum Mechanics by Bohr, Heisenberg, Schrödinger, Dirac and others.