

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: So we want to discuss line shifts and line broadening. And I assume you'll remember that last class we did some brainstorming and came up with quite a number of line broadening and line shifts mechanism, temporal lifetime broadening, motional broadening, external field broadening, collisional broadening. And as I indicated at the end of last class, those different broadening mechanism have much more in common than you may think.

And you will see that they have much more in common when we look at it from the fundamental perspective of coherence. Any line broadening mechanism comes because the atom experience, the environment and the drive field as a coherent source only for limited coherence time. So the concept of the coherence time will actually provide a common denominator for all those line broadening mechanisms.

However, before I present to you this kind of correlation function approach to what's line broadening, I think it's really important that we go through some simple cases. I always want you to learn phenomena in the simplest possible manifestation, in a situation where without any math you see what's going on. And what I particularly love is if there is an analytic solution where everything is sort of transparent here.

So therefore, before we discuss in a comprehensive way line shifts and line broadening, I want to go through simple cases. Now, I decided for this class that most of today's lecture uses pre-written slides because those cases are so simple that I would be almost afraid I would bore you if I would write it out. Because if you look at this equation in one second you get it, and it takes me 10 or 20 seconds to write it.

On the other hand, I would sort of give the responsibility to you. If I'm flooding you

with too much information, ask me questions, slow me down, or say, can you please go over that? So I don't want to obscure things by going through them faster.

I also looked a little bit forward to the end of the course. We are about OK with the pace of the course, or maybe one hour behind. And so in order to make room for discussions of superradiance, which I would like to have with you in a couple of weeks, I thought I can save half an hour here by going over that material a little bit faster. But these are really just illustrations of line broadening. And you should actually know most of them already.

So the first cases are actually leading you up to the Ramsey resonance, to the method of separated oscillatory fields, which is really absolutely important for high resolution spectroscopy. But before we can fully understand the Ramsey resonance, we have to understand the Rabi resonance. So what I've written down here is for you just the well-known formula for Rabi oscillations. And if you now simply assume we let the atoms interact with the field by a fixed amount time τ in such a way that we have a π pulse, then this probability as a function of detuning becomes a line shape.

And this line shape is plotted here. It has a full width of half maximum, which is approximately $1/\tau$. And this is-- you can regard it as the Fourier limit of finite interaction time broadening. We are observing the atoms for time τ and the line which is on the order of $1/\tau$. But we also know from the Rabi formula that we have these side lobes, so we see some oscillatory line shape.

OK, now, with the advent of trapped atoms, you can often have the situation that you have an ensemble of trapped atoms. You flash on your drive field for fixed amount time τ . At least for most of last century, this was not possible. You had atomic beams. And in an atomic beam, what is fixed is not the interaction time τ but the interaction length l . And then due to the velocity distribution, different velocity groups of atoms interact a different amount of time with that.

So for conceptual reasons but also for the historic context, you should sort of have an idea what the Rabi apparatus is. I may have mentioned it, but I regard Rabi as

sort of one of my ancestors in my family tree because Rabi was the PhD advisor of Norman Ramsey, who was the PhD advisor of Dan Kleppner, who was the PhD advisor of Dave Pritchard. And Dave Pritchard was my postdoctoral mentor. So I'm really talking here about my scientific great-great grandfather.

Anyway, so the famous Rabi apparatus is the following. You use one Stern-Gerlach magnet to prepare a certain hyperfine state. Then you have the interaction region, and this is what we are focusing on. And then later, if a spin flip has taken place, you can figure that out by running it through a Stern-Gerlach analyzer. So it's this ABC region, and we are talking about this middle region where we have an interaction time which is now given by the length divided by the velocity.

So what we have to do is-- well, that's why I'm saying it's simple cases, and I hope I can go fast-- you just take the previous result with a Rabi probability and convolute it with a velocity distribution. So when we do that, we find two effects. One is, well, due to the velocity distribution, the line width becomes a factor of 2 broader, but it's still proportional to something on the order of unity divided by the interaction time. And because the different velocity groups have a different kind of oscillations as a function of detuning, the oscillatory behavior now disappears because it becomes averaged out [? of ?] the velocity groups.

OK so this is the Rabi method or the Rabi resonance, where we have one interaction pulse or one interaction time for an atom. The next method is now what was introduced by Norman Ramsey as the method of separate oscillatory fields. And for that and other contributions he was awarded the Nobel Prize in 1990.

And the difference is the following. In the Rabi method, we let the atoms interact with the drive field for fixed lengths l or fixed time τ . But in the Ramsey method, we interrogate, we drive the atoms with two short pulses, which are separated.

So in the simplest case, we do 2π over 2 pulses separated by a time t . And now, if you have two pulses, everything is coherent. The two pulses can interfere constructively or destructively.

So what you would expect is that-- I want to give you two pictures. One is the Bloch sphere picture, but let's just sort of play with the concept of interference. If this is zero detuning, the two pulses are separated by a time t . And therefore, you would now observe, as a function of detuning, Ramsey fringes and oscillatory behavior, which has a spacing which is 2π over the time between pulses.

The envelope of the whole fringes is related to the short time of the duration. It's like a double slit experiment. You know, each slit is broad, and the width here is given by 1 over the short time Δt of the pulse. But then the two slits interfere. And the interference pattern is the distance between the slits, which now in the temporal domain is the time interval T between them.

So these are sort of typical Ramsey fringes. And now, if you would average over a broad velocity distribution, then you would kind of average-- you would maybe see one or two side lobes, but the other fringes are averaged out. So the central peak, the resolution is again on the order of $1/T$. It's the total time of the experiment, which is setting the ultimate resolution, similar to the Rabi resonance.

So this uses the picture of interference between two pulses. But I also want to sort of give you the Bloch sphere picture, because it's beautiful. For that I need a little bit of room.

So an atom enters the first Ramsey region, and it has to spin down. You do a $\pi/2$ pulse, which is at 90 degrees. And now what happens is there is a field-free region between the two pulses, so the atom is now precessing at its resonance frequency.

OK, but the synthesizer which is attached to your coils is also sort of precessing at this frequency. And when you add zero detuning, the synthesizer and the atoms are aligned again. And the second $\pi/2$ pulse is now flipping the atoms, and you've 100% excitation. And this is what you see here.

But now let's assume you are slightly detuned. Then the atom is precessing, and your synthesizer is precessing at a slightly different frequency. And now in the second region, the frequency of the synthesizer may be different with-- the phase of

the synthesizer is 180 degree different from the phase of the atom. And then instead of flipping the atoms up and getting 100% in the up state, the atom is now flipped down. And this explains the first minimum.

So you should really see that in this long region, nothing happens. But you accrue a relative phase between the synthesizer and the atom, which oscillates at its resonance frequency. So based on this model, you could work out mathematically every aspect of those fringes I've shown you. But I decided to take the equations out of the lecture and just present you with this physical picture. Questions?

AUDIENCE: If you have a velocity distribution, would the points of those minimums change or is it still this--

PROFESSOR: Well, good question. The spacing is the time T . But if you have a velocity distribution, let's say a velocity distribution which has a width Δv over v of 30%, what we have kept fixed in the beam experiment are the lengths, the interaction lengths. So time is length over velocity.

So therefore, when the velocity changes by plus/minus 30% in your atomic beam-- unless you take a supersonic beam with a very narrow velocity distribution-- then that means one velocity class has this set of fringes. The other velocity class has this similar set of fringes, but like in a harmonica, everything is now spread out by 30%. And if you have now 30% velocity resolution, that means you may be able to see two or three fringes.

But the central fringe is sort of like a white light fringe. There is no relative phase shift. And therefore all the different velocity classes will have a maximum at zero detuning. So in the extreme limit of a very broad velocity distribution, the only feature which survives is the central fringe. But this is where you obtain your spectroscopic information from. Will?

AUDIENCE: You have explained this Bloch sphere picture assuming unbroken coherence between your synthesizer and the field-free evolution of your atom?

PROFESSOR: Yes.

AUDIENCE: Is it equivalent to if you wrote down the time evolution operator, you would say that there's 0 Rabi frequency, but no detuning in your field?

PROFESSOR: Yes.

AUDIENCE: So what if I-- is this necessary, to have this unbroken coherence if I unplug my synthesizer and plug in a new one [INAUDIBLE]--

PROFESSOR: Yes, you can do spectroscopy with a resolution $\Delta\mu$, which is $1/T$, only if you have a synthesizer which has a frequency stability which is better than $1/T$.

AUDIENCE: OK.

PROFESSOR: Otherwise, we've done this a little bit when we talk narrowband and broadband. When we talked about narrowband and broadband cases, you are always limited in your resolution by whatever is broader. Here I'm discussing about the width of the atomic system, assuming a perfect synthesizer. But then in essence, you should convolute this result with the spectral distribution of your synthesizer.

And then if the synthesizer has a resolution which is worse than that, you would actually blur out the fringes through the convolution with the frequency spectrum of the synthesizer. [INAUDIBLE]?

AUDIENCE: So that also means if you add a [INAUDIBLE] to the second coil, [INAUDIBLE] just like if you had [INAUDIBLE] to one of the--

PROFESSOR: Yes, actually, if you would add a pi phase shift to the second coil, then the minimum-- the central feature would not be maximum or minimal. Yes. And that's why when you have atomic clocks with a beam, the question of distributed phase shifts within the microwave cavity play a big role. And this is related to are the two Ramsey zones really at the same phase or not? So that's an important issue for ultimately making resolution spectroscopy.

OK. Since-- OK. So we've discussed two methods now, the Ramsey method versus the Rabi method. And let me discuss advantages or disadvantages of the method.

So if you're an atomic physicist and you have to give advice to your friends whether they should use the Ramsey or the Rabi method, these are your talking points.

So one point is that I said the central feature of the Ramsey fringes is $1/T$. The Rabi feature is $1/T$, so we are both limited by time resolution, because this is a Fourier limit. But if you work out the details, you find that the Ramsey fringe is about 2 times narrower than the Rabi resonance. Just how things work out mathematically-- I can't give you any real deeper insight why. It just works out to be a factor of 2 narrower than the Rabi resonance.

There's one important aspect, and this is the following. When you accrue the spectroscopic information, you compare-- I hope you remember my demonstration-- you compare the atomic oscillator to the synthesizer, but you're not interacting with each other. So therefore you don't have any power broadening. You're comparing the free evolution of the atom with the propagating phase in your synthesizer.

And therefore, your Rabi signal has no power broadening at all. And that means that, however, the Rabi signal, at least for small detunings which are smaller than the Rabi frequency, will always depend quadratically on delta. Remember, generalized Rabi frequency, you add or make a Rabi squared and detuning squared in quadrature. So therefore you will always get effects for small detunings which are quadratic, whereas the way how you set up the Ramsey experiment, you can explore a linear dependence. So you have more sensitivity here.

Finally-- well, not finally. There are more. Next is the Ramsey spectroscopy. It's done in a field-free region, so you're not driving the system. You observe the free evolution. Therefore you have no AC Stark effects due to the drive field.

Of course, you drive the system in the two Ramsey zones. And some form of AC Stark effect may come in in the way that you may not have exactly a pi pulse due to AC Stark effects. But this is sort of a higher order effect. The basic spectroscopy is done by comparing the atomic oscillator with the frequency synthesizer.

Since it's a field-free region, this region can now be used for-- well, let me say

generally for experimental additions. And of course, what should immediately come to your mind is the Nobel-Prize-winning experiments of Serge Haroche where he had two Ramsey regions with microwaves. And in between in the field-free region, the atom passed through another cavity. And in between the Ramsey zones, the atom experienced a phase shift due to the presence of a single photon.

So the field-free region could now be used to put in a cavity which was filled with one or two photons. And the atom was in a non-destructive way reading out how many photons were there. So this is another advantage of the Ramsey spectroscopy, that you can now use the field-free region to measure something. You can introduce a phase shift, which can then be read out through the Ramsey interference.

Similarly, if you just think of how the resonance comes along, if you had a slightly fluctuating magnetic field between the two Ramsey regions, this would not necessarily broaden your signal. Because what you measure is the integrated phase evolution of the atomic oscillator. So in other words, what you get is shift of the resonance, which is the average over the inhomogeneous magnetic field. Whereas when you do a Rabi resonance, whenever you have a field inhomogeneity, you broaden and shift the resonance to this field value.

So in other words, the Ramsey resonance depends only on the average along the mean energy separation between the two levels. And therefore inhomogeneous fluctuations, spatial fluctuations, can completely average out. Whereas in the Rabi method, any kind of fluctuations leads to line broadening.

And finally, I will explain that in more detail later on. But what happens in the Ramsey method if the separation between the two regions is much longer than the spontaneous lifetime of one of the levels? Do you now get a resolution which is $1/T$, the temporal separation between the two interrogations? Or do you get a resolution which is $1/\tau$, the lifetime of the excited level? What do you think?

AUDIENCE: $1/\tau$.

PROFESSOR: It would be a good clicker question. So who thinks it's $1/\tau$? Who thinks it's $1/T$? So who thinks it's limited by lifetime? Who thinks it's limited by the interrogation time?

OK. So that's a minority for the Ramsey method. The minority is correct. And the picture you should have is that you remember the sort of picture of those oscillators.

But if some atoms decay away, it diminishes your signal. But the interference comes only from the survivors. And the survivors have survived in the exponential tail of the natural decay, but they are longer lived.

And therefore, you can actually do spectroscopy which is narrower than the natural line width, using Ramsey spectroscopy. But if you do Rabi spectroscopy, you're limited by the spontaneous lifetime. And this is probably what was in mind of the other people who raised their hand for option A. But why there is a difference between the Rabi and the Ramsey method, that's something I want to discuss later.

OK, so Ramsey has the possibility, sub-natural line width is possible when the interaction time is larger than the inverse natural line width. OK. Questions about Ramsey method? Good.

Physics Today has a wonderful article written by Ramsey which was reprinted recently, I think on the occasion of his death. I will post this article to our website. Then you can really read about it in the language of Norman Ramsey.

OK, let's move on. I said we are discussing simple examples. So we have discussed the example of Rabi resonance and Ramsey resonance. Now I want to talk about line shape with exponential decay.

One reason why I wanted to give you a simple model for exponential decay because in the end, everything is exponentially decaying because of the finite lifetime of levels. And with this very, very simple model, I want to convey to you that not all exponential decays are equal. You have to be a little bit careful. And this is just sort of the simplest example. And you learn something by figuring out what is different from spontaneous decay here and what are the consequences of that.

So remember where we are. We have the Rabi resonance. I gave you the simple example that the Rabi resonance is applied for fixed time τ . And then we did one extension we averaged over the velocity distribution.

But now we can just say, OK, we have our Rabi resonance here. But we assume that while we drive the atom, they decay away. If you want, you can think these are radioactive atoms and they decay [? radioactively. ?] For that situation, this model is exact. So therefore, instead of having a fixed interaction time τ , you have a distribution, which is an exponential distribution.

So all we have to do is we have to take our result, which I discussed 10 or 20 minutes ago, with the fixed interaction time and convolute it with the distribution of times the atom experiences the drive field. And what I've introduced here is the exponential is γ . And the mean interaction time over which an atom interacts with the drive field is just $1/\gamma$, and I called that τ .

So now we had the Rabi probability here, but now we convolute it with the distribution of interaction times. And this is now the probability that after the Rabi pulse the spin was moved from spin down to spin up. Or if you have an electronic transition, from ground to excited state.

So this integral can be analytically solved. That's why it's worth presenting. And what you get is a Lorentzian line shape. And this Lorentzian line shape shows power broadening, which actually you should find nice, because we will sometimes [INAUDIBLE] in power broadening you can't get out of perturbation theory.

And a lot what we have done and what I actually want to do for the remainder of this chapter on line broadening is a perturbative approach. So that's another reason I want to present it to you here. These are some non-perturbative results, and they show the physics of power broadening, saturation broadening.

But there are two things which are noteworthy. One is the full width at half maximum is not γ but 2γ . So if we had natural decay at a rate γ , the Lorentzian which we get is only half as wide. But you can immediately say, well, that

can be understood because here I assumed the atoms just decay away no matter whether they are in the ground or in the excited state.

And I gave you a model that you assume there is maybe radioactive decay independent of the internal state. And now you can wave your hands and say, OK, if you have only decay in the excited state and not in the ground state, this should give you a factor of 2 in the width. And this may explain why we have a full width of half maximum here of 2γ and in spontaneous decay it is γ .

But there's another thing which is interesting, maybe more interesting. And this is the power broadening. If you take the power-broadened Lorentzian line width and we look at it in the limit of high power, it is 2 times the Rabi frequency. Well, if we have a system which has spontaneous decay and we would go to the high power limit-- we've discussed it before. What you get is square root 2 of the Rabi frequency.

So the message I can give you here is that saturation broadening, power broadening depends sensitively on the exact nature of the decay and of the lifetime broadening involved. And if you really want to do it right, you have to use the optical Bloch equations.

So let me just write that down. So what we learned from that is, yes, we get power broadening. We have a simple model for power broadening here. But power broadening depends sensitively on the nature of the decay process. And so if you want to get this result without any assumptions or approximations, you should use the optical Bloch equations.

And yes, your homework assignment looked at the optical Bloch equations. And I think you also found out that some results for the optical Bloch equations really depend on the ratio of γ_1 γ_2 or T_1 and T_2 . So that is with a more mathematical formalism shows you that the way how you introduce decay into the atomic system, it's not just there is one time constant and the result just depends on the time constant. There are some subtleties. Questions about that?

OK. Now-- and many people have asked me about it-- I think for the first time in this course we bring in motion of the atoms. So the atoms are now not pinned down at the origin. Maybe you can imagine you have an atom which is held in a solid state lattice with a nanometer, and we just look at the internal structure. We have maybe some ions implanted into a material. And these ions are fluorescent, and we are probing them. Or we're doing spin flip on nuclei, which are nuclei of atoms which are part of a condensed metal lattice.

Or if you want to think more [? with ?] the methods of atomic physics, you have the most tightly confining ion trap in the world. You are deeply in the Lamb-Dicke limit, and your ion just cannot move. It's always in the ground state of the ion trap, and all you are dealing with is the internal degree of freedom.

Actually, let me make a comment. I often see when people approach me and ask me question that they are not necessarily making the separation. When they think about what happens to the internal structure of atoms, what is in their head is, but there is motion, there is recoil. You can, in my experience, always separate the two. You can create a situation where you only probe the internal degree of freedom by tightly confining the atom, and then you can relax the condition that the atoms is tightly confined. Now the atoms can move, and then all of the things we want to discuss now come into play.

Sometimes people assume, yes, but if you confine an atom, doesn't the atom always have a de Broglie wavelengths? And isn't that another length scale? The answer is no, because you need a coupling between the internal degree and the external degree, or you need some way of exciting the external degree of motion. And if you have a tightly confining ion trap, $\hbar \omega$, the next vibration or level in the ion trap is so high, you may not excite it with the recoil of a photon.

But there is another limit which I often find very useful, and this is the following. When we talk about spectroscopy, spin flips, electronic transitions, we have not really talked about the mass of the atom. The mass of the nucleus only appeared in the reduced mass. Remember when we did last class the hydrogen atom? We had

the reduced mass, which was slightly different from the electron mass.

So if you want to completely exclude the motion of the atom, just work in the limit that the nucleus has infinite mass. If the nucleus has infinite mass, its de Broglie wavelength is 0. It's confined in a harmonic oscillator to sort of a delta function. So by just assuming that you deal with infinite mass, you automatically neglect all possible motions.

And as you will actually see from the next formula, working in the infinite mass limit means that your Doppler shift is 0, your recoil shift is 0, everything is 0. So either way, I've given you now two ways how I recommend that you think about all the physics we have discussed which deals with the internal degree of freedom by either saying the atom is tightly organized. But then some people say, ah, but then it has Heisenberg Uncertainty. There's a lot of momentum.

It doesn't matter. If it's localized, it's localized. But you can also assume just the infinite mass limit. And in both cases, the result is you can completely talk about so many aspects of internal excitations without even considering what happens externally. Any questions about that?

So now we take the mass from infinity to finite value, and now we want to allow the atom to move and have kinetic energy. So let's start out very simple. We have an atom which is addressed in the excited state, and it emits a photon.

Before the emission of the photon, the total energy is the excitation energy. But after the emission of the photon, the atom is in the ground state. The photon has been emitted.

But now the atom, due to the photon recoil $\hbar k$, has kinetic energy. So therefore the frequency of emission of the photon does not happen at the resonance frequency of the atom because some part of the electronic energy goes into the kinetic energy of the recoiling atom. And this is called the recoil shift. We start at 0 velocity. At 0 velocity, you don't have any Doppler shift. But you do have recoil shift.

Well, we can play the same game. We have an atom addressed. It absorbs the

photon. And after the atom has absorbed the photon, it's in the excited state.

But now you have to excite the atom if you want to transfer the atom to the excited state with a frequency which is slightly higher than the resonance frequency. Because the laser has to provide not only the energy for the electronic excitation but also the energy for the kinetic energy at which the atom recoils. So therefore we find that, due to the recoil of the photon, the absorption line and the emission line for an atom addressed are shifted. The shift is the recoil energy, $\hbar^2 k^2 / 2m$ -- the momentum of the photon squared, divided by 2 times the mass of the atom. And the shifts are opposite for absorption and emission.

So therefore, if you look at the two processes for absorption and emission, there is a recoil splitting between the two. This recoil splitting between emission and absorption is just a few kilohertz. And it was really one of the wonderful accomplishments when high resolution spectroscopy came along and John Hall at Boulder managed to have lasers stabilized to sub kilohertz. For the first time, this recoil splitting could be resolved. So he had set up some intracavity absorption, and he saw sort of two peaks in some kind of spectrum. I don't remember the details, but two peaks split by a few kilohertz were really the hallmark of the photon recoil shifting the lines away from resonance.

OK, so now we know what the kinetic energy of the atom does. If it emits a photon, there is recoil. But now, in addition, we can drop the assumption that the atom is initially at rest when it absorbs or emits. Now the atom is moving.

But for that, we don't need any new concept. Because the moving atom-- we can just do a transformation into the frame of the atom where the atom is at rest. And then just using the relativistic transformation, we are now transferring the laser frequency from the atomic frame into the lab frame.

So what I've written down here is the general special relativity formula for the frequency shift. And I've assumed that the photon is emitted at an angle ϕ with respect to the motion of the atom. So now we obtain-- OK, let me do a second-order expansion. Usually, our atoms are non-relativistic, so it's the first- and second-order

term which are most important. And if we are now looking in the lab frame, what is the frequency where we emit and absorb photons. It's a resonance frequency minus/plus the recoil shift, which we have already discussed in isolation by assuming we have atoms at 0 velocity.

But now the velocity of the atom leads to a first-order and second-order Doppler shift. If v/c is small, what is most important is the first-order Doppler shift. And in almost all cases [? where we ?] do spectroscopy, dominant line broadening effect comes from the first-order shift.

However, let me point out that the first-order shift can be suppressed. One simple way to suppress it-- interrogate the atoms at an angle of 90 degrees. Have an atomic beam. And if you interrogate them with a laser beam at 90 degrees, the cosine ϕ is 0. $k \cdot v$ is 0. And this is the oldest method to do Doppler-free spectroscopy.

In your new homework assignment, you will discuss saturation spectroscopy. If you have a broad velocity distribution, but you find a way of labeling atoms with a certain velocity class, then you have created your own narrow velocity class where the Doppler broadening is absent because you've only one velocity class. And these methods of nonlinear spectroscopy, the concept will be developed in the homework assignment.

Finally-- and this will be the next chapter which we talk in class here in about two weeks-- this is by having two-photon spectroscopy. To give you the appetizer, if you have two photons from opposite direction, the Doppler shifts. One has a positive, one has a negative Doppler shift. And if you stack up the two photons, the sum of the two Doppler shifts is 0. So two-photon spectroscopy provides you an opportunity to completely eliminate the Doppler shift.

However, no matter what you pick here for the angle, there is a part of the second-order Doppler shift you can never get rid of. And this is something important to keep in mind. The second-order Doppler shift, at least one part of it, comes from the relativistic transformation of time. So if you have atoms moving at different

velocities, time in the frame of the atom ticks slightly differently depending what the velocity is.

And therefore, if you do spectroscopy, you measure time in the lab frame, but the atoms measure time in the resonance frequency in their own frame. Then there is inevitably broadening. So for instance, when people did Ted Hansch's experiment, the famous two-photon spectroscopy on hydrogen, high resolution, determination of the Lamb shift, the Rydberg constant-- one of the flagship experiments of high resolution laser spectroscopy.

A limit is the second-order Doppler effect because of its relativistic nature. And the only way to suppress the relativistic Doppler effect is by cooling the atoms, reducing their velocity. So let me just write that down. So when we suppress the first-order effect, then the limit is given by the second-order term.

And just repeat, the second-order term cannot be eliminated by playing geometric tricks, 90-degree angles and such, because it's fundamentally rooted in the relativistic definition of time. OK. Any questions about recoil shifts, Doppler shifts? Yes.

AUDIENCE: What are the correct way to determine whether we're in tightly confined [? regime ?] or not?

PROFESSOR: OK, the question is what is the criteria on whether we are tightly confined or not. I will give a full answer to that in about a week when I discuss with you in great detail the spectrum of a confined particle. And what we will introduce is it's the frequency of harmonic confinement. And we have to compare the frequency of harmonic confinement to two other relative, to two other important frequencies.

One frequency is the recoil frequency. And another one may be the natural line width. So in other words, when we discuss the spectrum of confined particles, we can discuss it as a function of three parameters-- confinement frequency, recoil frequency, and natural line widths. And based on the hierarchy of those three frequencies, we find limiting cases.

And we will find, then-- and this is probably what you are aiming for-- at some limit when you reduce the confinement in your harmonic oscillator, you will actually retrieve the free gas limit. Or to be very brief, confinement, you have the benefit of confinement. Confinement means the motion of the atom is quantized in units of $\hbar\omega$. So you shouldn't think about velocity. You should think about discrete levels.

And I will show you that the spectrum which is broadened becomes a spectrum with discrete levels and side bands. As long as you can resolve the side bands, you can see them, you can actually address the line in the middle which has no motion blurring at all. But you have to resolve it.

So one condition here now is that the harmonic oscillator frequency is larger than the natural broadening of each line. If the lines blur, you're pretty much back to free space and you've lost the advantage of confinement. But maybe we can discuss some of those aspects after I've introduced the line shape of confined particles.

OK. OK. Let me now discuss briefly, or use what we have just discussed about Doppler shift to discuss the line shape in a gas. Well, of course, line shapes in a gas-- that's what all people observed when they did spectroscopy before the advent of trapping and cooling. But even now, we often have a situation that we have a thermal clouded microkelvin temperature. And what we see is still broadening due to the thermal motion.

So therefore, let me just tell you a few aspects of that which you might find interesting. So one is when we have a non-degenerate gas, this is described by a Boltzmann distribution, sort of a Gaussian distribution of velocity. And therefore, if the first-order Doppler shift is relevant, the first-order Doppler shift is proportionate to v , so therefore the spectrum we observe is nothing else than the spectrum in velocity multiplied with a k vector.

The Doppler shift is $k \cdot v$, and therefore the velocity distribution by multiplying it with k is turned into a frequency distribution. And so the classic frequency

distribution you would expect in a gas is simply the Gaussian distribution. And the Doppler width, the spectroscopic width, is nothing else than whatever the characteristic speed in your Boltzmann distribution is, typically the most probable speed, 2 times the temperature over the mass multiplied with the k vector.

And in many cases, it is this Doppler width which dominates. I've just given you typical examples here. If you've stabilized your laser to a room temperature vapor cell, you will encounter typical Doppler widths on the order of 1 gigahertz. This is 100 times larger than the natural line width, which is 100 times larger than the recoil shift. So that is the usual hierarchy of shifts and broadening mechanisms.

So therefore, you would think that if the Gaussian line widths due to the velocity distribution is 100 times larger than the natural line widths which is described by Lorentzian, you can completely neglect the Lorentzian. But that's not the case, and that's what I want to discuss now.

What happens is you have a Gaussian which is much, much broader. But the Gaussian decays exponentially, whereas your narrow Lorentzian decays with a power law. So just to give you the example, if you go two full half-line widths away from the center of a Gaussian, the Gaussian has dropped to 0.2%. The Lorentzian has still 6%.

So therefore, what happens is if you have your gigahertz broadened line in a gas, but you go further and further away, at some point what you encounter are not the Gaussian wings but the Lorentzian wings. And that's maybe also sort of intellectually interesting. I've talked to you about homogeneous, inhomogeneous line widths. The bulk part of the Gaussian is inhomogeneous because you can talk to different atoms at different velocities. Each atom resonates with its own Lorentzian, and it is inhomogeneously broadened.

But in the line widths which is due to the Lorentzian, the homogeneous broadening dominates. And since the Lorentzian has information about either the natural line widths or in a gas about collisional broadening, you can actually, far away in the wings of your line shape, retrieve information about collisional physics, which, in the

center of the line, is completely masked by the first-order Doppler shift. OK. So I've already written it down. So far wing spectroscopy in the gas phase can give you interesting information about atomic collisions and atomic interactions.

So having started out by telling you that it's a first-order Doppler shift which usually dominates, but then telling you that if you go far away from the line center, the Lorentzian actually dominates, that means now there are situations where we want actually both. And the general solution, of course, is you do a convolution. Each atom with a given velocity has a Lorentzian. And then you have to do the convolution with the velocity distribution.

So therefore, the general situation for gas phase spectroscopy is the convolution of the Lorentzian for each atom and the Gaussian velocity distribution. And since this was the standard case which people encountered when they did spectroscopy in the gas phase, this convoluted line shape has its own name. It's called the Voigt profile. Colin.

AUDIENCE: It's not obvious to me what exactly about the atomic collisions and interaction, the wings of the Lorentzian [INAUDIBLE]. Don't you just learn about the bare line? Like the un-Doppler shifted line itself?

PROFESSOR: OK, that's a good question, as it is not obvious. And yes, the literature is full of it. Because if you don't have Doppler-free spectroscopy, if you are always limited by the Doppler broadening, put yourself back into the last century. But you really want to learn something about how atoms interact and collide, this was the way to do it.

I don't want to go into many details because it's a little bit old fashioned. We have other ways to do it now. But I find it intellectually interesting when we talk about line shapes to realize there are maybe some subtleties we wouldn't have thought about it by ourselves. But just to give you one example, for whatever reason, you're very, very interested. What is the rate of collision of atoms?

Let's just assume the simplest model for collisions that when two atoms collide, the excited state is de-excited. Then the lifetime of the excited state is no longer $1/\nu$

gamma, the lifetime is the collision time defined as a de-excitation time. And you're really interested for whatever reason, because you have the world's best theory on this object, that you have a theory what is the collision time for excited sodium with argon, with helium, kind of with all different elements. And you have an interesting theory which actually reflects how sodium in the excited state would interact with noble gases. And you really want to test your theory.

Well, now what happens is the situation is simple. You have a Lorentzian, but the Lorentzian has a broadening which is the collision rate. And by carefully analyzing the wings of your Doppler profile, you find the collision rate-- as a functional of buffer gas density, as a function of noble gas, whatever you pick. That's one example.

Another example would be, there may be also somewhere nontrivial shifts when the atoms collide. We'll talk about it also a little bit later. The experience-- de-excitation is one possibility. But more subtle things can happen. For instance, just a phase perturbation, or when atoms come close to each other, they feel the electric field. And the electric field causes an AC Stark shift.

But by understanding the AC Stark shift or AC Stark broadening which comes along with that, you can maybe map out the interaction potential between two atoms. So people were really ingenious in trying to learn something about atomic interactions from spectroscopic information, and that was one of the few tools they had. Other questions?

OK. We have covered the simple examples. And now I want to give you a more comprehensive framework called perturbation theory of spectral broadening. And in the last class, I mentioned to you already that, by using this framework, we can more deeply understand line broadening mechanisms. And one highlight actually will be that next week, using these concepts, we will actually understand that collisions cannot only cause collisional broadening, they can also cause collisional narrowing.

So some things which are counterintuitive have a very clear description using this

method. So I think what I'm able to do in the remaining 20 minutes is to step you through the derivation. And then we apply it to a number of interesting physical situations next week.

Now, in a way, I have to apologize that what I present to you is time dependent perturbation theory again. And again with a slightly different notation, so I think I go rather quickly for the part which is just review. But repeating something is also a good thing. But then I will tell you when we go beyond what we have discussed and beyond what you may have seen in textbooks.

So in other words, we do time dependent perturbation theory. We have a wave function which is in two states, A and B. And there's a time dependent perturbation v . Schrodinger's equation tells us that it's probably the interaction picture, that the rate of change of the amplitude B comes because we start in the state A and the perturbation couples from state A to B.

And so we can solve it. At this point, it's not even perturbative, it's general. And we are interested in spectroscopy in the rate of a transition, because we do spectroscopy and we measure what is the population in the excited state. Because there was a rate at which atoms were transferred from the calm to the excited state. So the rate is the probability to be in the excited state per unit time.

So what we are interested now is, what is the amplitude B? And what is the probability B squared to be in the excited state? And now comes perturbation theory. If we take Schrodinger's equation and we integrate it with respect to time-- so we go from \dot{B} to B-- we are integrating here with respect to time. And in first-order perturbation theory, we assume the initial state is undepleted. And we replace the amplitude A of t by its value at time t equals 0, which is assumed to be unity.

OK. So with that-- oh yeah, and this may be something. I'm not doing anything which goes beyond perturbation theory, but I'm using a slightly different formulation. Because you will see that I need it in a moment. The rate is $d(B^2)/dt$. And if I take the derivative of B squared, I get $2B \dot{B}$.

So now I'm using the perturbative-- I'm inserting this function B into this expression for the rate. And this is what I obtain. So I can take the time derivative of it, but the time derivative is only affecting the upper integral.

So therefore what I get is B is the integral. B dot is the integrand. And now I get this expression, which has the product of the two matrix elements.

Simple mathematics, plug and play. No new concept. But what it leads us now is, and this is what is usually not so much emphasized in perturbation theory, that everything which happens to the atom, and this is the rate at which we excite the atom, is now involving a correlation function. It involves sort of an integral over the drive field v at time t and time t' .

And this is the important concept when you want to explain and understand line broadening and such. You are driving the system with an external field. And often in perturbation theory, you assume the external field is just $e^{i\omega t}$, and this correlation function is just $e^{i\omega t}$. And it's so trivial that you don't even recognize that $e^{i\omega t}$ is not the time dependence of your field. $e^{i\omega t}$ is sort of the product of the field at $t = 0$ and the field at time t and the field at time t' .

But if you have a more general field with lots of Fourier components, the difference between whether it's a correlation function or the field itself becomes important. In other words, I'm now telling you whenever you did perturbation theory, this is what you did. Maybe you didn't notice it, but what you had was actually the correlation function between the drive field at two different times, t and t' .

OK, so our rate is now given by the correlation function of the field. And then we so-to-speak Fourier transform it with $e^{i\omega_0}$.

OK. Let's just streamline the expressions, make them look nicer. We integrate between time 0 and t . But let's now assume-- which is actually the situation for many fields of interest-- that the field is invariant against translation in time t . So therefore this correlation function does not depend on two times. It only depends on the time

difference τ .

Finally, because of the complex character of the Schrodinger equation, I had an expression but I had to add the complex conjugate. Remember, the rate was the derivative of B^2 . And the derivative of B^2 is $B^* \dot{B} + \dot{B}^* B$.

You get two terms. And this is carried forward with the complex conjugate. But if the correlation function has the proper t , that complex conjugation means you can go to negative time. $e^{i\omega t}$ complex conjugate is $e^{-i\omega t}$. That means now that we can absorb the complex conjugate by integrating not from 0 to t , but having the integral from $-t$ to t .

And this will be the next step for most situations of interest. This correlation, if you drive it with the field, the field has a finite coherence time. So therefore this integral will not have any contribution when the times are longer than the coherence time. And then we can set $-t$ and t to infinity. So that will be our final expression which we will use to discuss line broadening and line shifts.

But we are not yet there. We need-- so far, I've just done ordinary perturbation theory. The one extra thing is I'm stressing that when we have a product of-- when we had a matrix element squared in perturbation theory, this is really a correlation function between the [external, ?] the drive field at two different times.

We come back to that when I discuss the result. But the second thing I want to introduce now is that this framework, which I have formulated, allows me now to include the fact that different atoms in my ensemble may experience a different drive field. For instance, I gave you the example in last class, if you have Doppler broadening, you have atoms which start out at the same point. But the faster ones move faster and experience the laser field now with a different phase.

So different atoms now experience the perturbation v in a different way. So what I've done here so far is I've pretty much written down Schrodinger's equation for a single particle. But now we have to do an ensemble average.

So therefore I introduce now an ensemble average by just taking that expression and averaging over all atoms in the ensemble. So then I get the ensemble averaged rate. All the correlation functions we discussed, our ensemble averaged correlation functions and our final result will also have an ensemble average.

OK. So this correlation function between v of 0 and v of t will go to 0 for very long times. Even the most expensive laser in the world, the electric field which is emitted now is not related to the electric field which is emitted in an hour. Because the phase relationship has been lost, and therefore the correlation function has decayed to 0.

OK. So this is the ensemble average. So therefore, what I'm naturally drawn to now is that if I take this correlation function, and I know any correlation function has a characteristic time called the coherence time where it decays. And therefore I can now discuss two limiting cases. One is where the time evolution of the system is started for times much shorter than the coherence time or much longer than the coherence time.

And if what I'm telling you right now reminds you of my discussion of Rabi oscillation versus Fermi's golden rule-- yes, this is actually a very analogous discussion. OK, so there are the two limiting cases. If the time is much shorter than the coherence time-- let me give you the example of an oscillating single mode field. The perturbation v of t is just oscillating with one frequency, ω . And that means if I look at the correlation function at time t and time t plus τ , it is simply the amplitude of the field squared times e to the $i \omega \tau$.

And now I can take this correlation function; put it into my integral, which has just disappeared from the screen; do the integration with e to the $i \omega_0 t$. And this is the result I obtain. And of course, this is nothing else than what you have always obtained in time dependent perturbation theory with a sinusoidal field. It is this characteristic sine detuning t over detuning, which in the limit and in the limit of when you square it and go to the limit of long times, it turns into a delta function. This gives us Fermi's golden rule.

And of course, it has the same behavior at short times. At short times, the probability for the atom to be in the excited state is quadratic. Quadratic is like an incipient Rabi oscillation. And in perturbation theory, we never get higher up. We just look at the beginning Rabi oscillation.

So therefore, the probability is quadratic. But I'm talking about the rate, and the rate is probability divided by time. So that means the rate is linearly increasing in time. So I'm just saying this is nothing else than rewriting the physics of Rabi oscillations.

OK. If the time is longer than the coherence time, then we integrate the integral, not from minus t to plus t . We can take the limits to infinity. And that means we obtain a result which is now independent of time. And that means since W_{ba} is the rate, we retrieve a constant rate, and this is what we have done in Fermi's golden rule.

So therefore, when we look at the time evolution of a system driven by an external field, the moment we look for the time evolution longer than the coherence time-- and this is where the main interest is in spectroscopy-- we have a Fermi's golden rule result that the system is excited at a constant rate. And I want to now reinterpret this rate.

This rate is nothing else than the Fourier transform of a correlation function. It makes a lot of sense. You apply time-dependent magnetic fields, perturbations, fluctuating magnetic fields, whatever, vibration and noise in your lab. You just apply that to the atom, and the atom is nothing else than a little Fourier analyzer. It says, my resonance frequency is ω_0 . And all of what matters for me to make a real transition is what you offer me at ω_0 .

And I now Fourier analyze whatever acts on me, the correlation function of the perturbation which acts on me, and I fully analyze it. And what matters for my rate to go to the excited state is the Fourier component at the resonance frequency. It's just a generalization of what we have done in perturbation theory when we assumed that we have a drive field only at one frequency.

So I've written it down here for you. The rate of excitation is nothing else than the

Fourier transform of this correlation function. But let me now take it one step further, which also makes a lot of sense. The Fourier transform of the correlation function-- the correlation function is the convolution of the time-dependent fields with themselves, v of t with v of t plus τ .

The Fourier transform of the convolution is the product of the Fourier transform of the field itself. So therefore, I can take whatever perturbation the atom experiences in its own frame-- external fields, moving around. Whatever the atom is exposed to, I have to calculate the power spectrum of what the atom feels. And this power spectrum provides me the excitation rate. It's nothing else than Fermi's golden rule but generalized to the concept of an arbitrary spectrum of the driving field.

Questions so far?

You can also say that's a wonderful way to look at energy conservation. If an atom is exposed to any kind of environment, it goes from the ground to the excited state only to the extent that whatever acts on the atom has a Fourier component at the resonance frequency. And it is only the power of the fluctuating drive field at the resonance frequency which is responsible for driving the atom.

And this is energy conservation. The frequency component has to be ω_0 to take the atom from the calm to the excited state. All the other frequency components take the atom to the virtual state and take it down again. They create maybe line shifts or something like this. But a real transition, a transition where the atom stays in the excited state, requires photons at the resonance frequency. And so to speak, this measures how many photons are acting on the atoms.

Let me now give you one or two general features of such correlation functions which I just find very, very useful. And then I think your time is over. If we have G of ω is now the spectrum of the fluctuating fields. And let's assume, yes, eventually we have a fluctuating field which is somewhere centered at the resonance frequency.

After all, we use a laser, but the atom may now move around in the laser beam. The mirrors may be vibrating. So the spectrum seen by the atom is sort of broadened around the resonance frequency. And the broadening is nothing else than 1 over

the coherence time of the environment.

So let me just normalize the correlation function that the integral is unity. And then this is trivial but important. The value at the resonance frequency is 1 over the broadening and is therefore the coherence time. It's sort of subtle but important. If you have a normalized spectrum, the more coherent your source is, the larger is the value of the correlation function in the center.

Let me do the Fourier transform. If you Fourier transform something like this, it gives an oscillating function at the resonance frequency. But let me just multiply by $e^{-i\omega t}$ and sort of shift everything to 0 frequency, then we would find that the temporal correlation function decays. And it decays over characteristic time τ_c coherence, which is nothing else than the inverse line width of the Fourier transform. This has nothing to do with atomic physics. It's just properties of the function is Fourier transform.

But now we have the situation. Our rate was the temporal correlation function times $e^{-i\omega_0 t}$. The integrand here is exactly what I'm plotting here. And so if I perform the integration, at least without getting the last numerical factor, I can approach the result by the correlation function at time t equals 0 times-- if I do the integration-- times the temporal rates of this curve, so this is the correlation time τ_c .

So therefore, if I have, for instance, my operator is the electric field and I drive the atom with a dipole operator, what I find is the correlation function at t equals 0 is nothing else than the electric field squared, which is what we have called the Rabi frequency squared so far. But now I multiply it with a coherence time.

So this result should come very naturally to you because when we have Fermi's golden rule, we have a matrix element squared times the delta function. But I've emphasized it again and again-- the delta function is representative for spectral widths for density of states. And if we have an environment which causes spectral broadening, 1 over the coherence time is nothing else than the spectral widths here.

And so I've done here exactly what the delta function Fermi's golden rule asked me to do. So if I had wanted, I could have just said, look, here is Fermi's golden rule. And by interpreting Fermi's golden rule the way I just did, I could have written down this result for you right away.

OK, I think time is over. And maybe the summary which I could give you now is a good starting point for our lecture on Monday. Reminder, Friday we have the midterm in this other lecture hall in the other building. We start at the normal class time. Please be there on time.