Vol. 2

Chapter 2.2

Page 1:

Alright everyone, welcome back to Physics 608. I'm Distinguished Professor Dr M A Gondal, and today, we are embarking on a new and truly central topic in laser spectroscopy, which you'll find in your notes as Chapter 2.2.

We're going to build up, from first principles, one of the most powerful techniques we have for achieving ultra-high spectral resolution.

Page 2:

So, the title of this chapter is the "Saturation of Inhomogeneous Line Profiles." Now, this might sound a bit abstract at first, but what we're really talking about is a suite of incredibly clever techniques that allow us to overcome one of the biggest obstacles in precision gas-phase spectroscopy.

So let's motivate this. Why are we dedicating an entire section to this topic? What problem are we trying to solve? The short answer, as we'll see, is the Doppler effect. In any real-world experiment with atoms or molecules in a gas, they aren't sitting still. They're whizzing around with a wide range of velocities, and this thermal motion smears out the beautiful, sharp spectral lines that quantum mechanics predicts. Saturation spectroscopy gives us a toolkit to circumvent this smearing, to look past the crowd and interrogate individual groups of atoms. This allows us to reveal the true, underlying structure of our transitions. We're going to develop the quantitative theory behind this, and by the end, you will understand exactly how we can use lasers to burn holes in atomic velocity distributions and observe the famous "Lamb dip," which is a cornerstone of modern metrology. Let's begin.

Page 3:

So, let's lay out the core motivation and the phenomena we'll be studying.

First, as I just mentioned, high-resolution laser spectroscopy in gases is almost always limited by Doppler broadening. Think about it: you have a gas of atoms at room temperature. They're moving with speeds of hundreds of meters per second. This thermal motion causes the absorption and emission frequencies to be shifted up or down depending on whether the atom is moving towards or away from your laser. The result is that a spectral line that should be, say, a few megahertz wide, which we call the natural or homogeneous linewidth, gets smeared out into a massive feature that can be a gigahertz wide or more. This completely obscures any fine or hyperfine structure we might want to resolve. It’s like trying to read fine print through a blurry lens.

This brings us to the second point. Saturation techniques are our solution. They provide a method to, as the slide says, "dig below" the Doppler envelope. This is a fantastic analogy. The Doppler profile is like a huge, broad hill, and hidden underneath are sharp, narrow features. Saturation techniques use the intensity of the laser itself as a tool. By using a sufficiently strong, monochromatic laser, we can selectively interact with only a very small subset of the atoms—those with just the right velocity to be resonant with our laser. By altering the population of this specific velocity subgroup, we create a spectral marker that is not subject to the full Doppler broadening of the entire ensemble.

And what are the key signatures of these techniques? We'll encounter two main phenomena, which are really two sides of the same coin. First, we'll talk about "Bennet holes" or "burned holes." This refers to the literal depletion of ground-state atoms within a specific velocity class, creating a 'hole' in the Maxwell-Boltzmann velocity distribution. Second, and perhaps more famously, we will derive the "Lamb dip." This is a sharp, sub-Doppler dip that appears right at the center of a Doppler-broadened absorption line under the right conditions. This dip becomes an incredibly precise frequency marker.

Page 4:

So, what is our goal for this section, and what is the ultimate payoff?

Our primary goal is to be rigorous. We will derive, step-by-step, the quantitative expressions for these effects. We're not going to rely on hand-waving arguments. We will start with the fundamental light-matter interaction, incorporate atomic velocities, and build up the mathematical framework to predict the exact shape, width, and depth of these saturation features. I want you to understand what every symbol in every equation means, and every assumption we make along the way. By the end, you should be able to look at a saturation spectroscopy experiment and model it from the ground up.

And this is not just an academic exercise. The practical pay-off of understanding this material is immense. These techniques are at the heart of many cutting-edge applications. For example, frequency locking of lasers. That sharp Lamb dip provides a perfect error signal to lock a laser's frequency to an atomic transition with sub-megahertz accuracy, or even much better. This leads directly to precision metrology, like atomic clocks, and even the modern definition of the meter, which is based on the frequency of a Lamb-dip-stabilized laser. These methods can also be used for things like isotope separation, where you selectively excite the isotope of interest. And finally, by measuring the positions of these ultra-narrow spectral features with incredible precision, we can perform some of the most stringent tests of fundamental physics, like measuring the fine-structure constant or searching for tiny effects of parity violation in atoms.

So, the physics we're about to unpack is not just beautiful, it's incredibly useful. Let's start by reviewing the two fundamental types of line broadening.

Page 5:

Alright, let's begin with a quick review of the two principal broadening mechanisms. It is absolutely essential to have a crystal-clear understanding of the distinction between homogeneous and inhomogeneous broadening before we can tackle saturation.

First up is homogeneous broadening. The key concept here is that the broadening mechanism affects every single absorber in the ensemble in exactly the same way. Every atom has the same center frequency and the same linewidth. If you were to somehow pull out a single atom and measure its absorption profile, it would have the same width as every other atom. Any given atom has a chance of absorbing a photon anywhere within this homogeneously broadened line profile.

What causes this? The dominant sources are processes that limit the coherent lifetime of the quantum states. This includes spontaneous decay—the natural lifetime of the excited state, dictated by quantum electrodynamics. It includes collisions, which interrupt the phase of the atomic oscillator. And, critically for our upcoming discussion, it includes power broadening, where a very intense laser field itself drives transitions so rapidly that it effectively shortens the lifetime of the states, broadening the line.

The characteristic lineshape for all these processes is a Lorentzian. This mathematical shape arises naturally from the Fourier transform of an exponential decay, which is the hallmark of these lifetime-limiting processes.

And the final bullet point here is the most important defining feature: every particle, every atom or molecule, has the exact same identical resonance frequency, ω 0 ω_0.

Page 6:

Now for the second category: inhomogeneous broadening.

In this case, the width arises from a statistical distribution of resonance frequencies across the entire ensemble of atoms. This is a crucial difference. Individual atoms in the sample do *not* all have the same resonance frequency. Each atom still has its own, very narrow, homogeneous linewidth—its personal Lorentzian profile—but the center frequencies of these Lorentzians are spread out over a wide range. The overall lineshape we observe is the envelope of all these shifted individual profiles.

The classic, and for us, the most important example, is the Doppler shift resulting from molecular thermal velocities. As we discussed, atoms moving towards the laser see an up-shifted frequency, and those moving away see a down-shifted frequency. Since the velocities of the atoms follow a Maxwell-Boltzmann distribution, the resulting distribution of resonance frequencies is a Gaussian. Therefore, the characteristic lineshape for Doppler broadening is a Gaussian.

So what happens in a real gas, where you have both types of broadening? Well, both mechanisms act simultaneously. Each individual atom has its own homogeneous Lorentzian profile, and the collection of all atoms has an inhomogeneous Gaussian distribution of center frequencies. The observable lineshape is therefore the convolution of the Lorentzian and the Gaussian. This resulting lineshape has a special name: the Voigt profile. At its core, it's Gaussian, but in the far wings, it has the characteristic 1 ω 2 $\frac{1}{\omega^2}$ fall-off of a Lorentzian.

Page 7:

This brings us to the central point of this whole chapter. A saturating radiation field—that is, an intense, monochromatic laser—interacts profoundly differently with these two types of broadening.

For a homogeneously broadened line, where every atom is identical, an intense laser "talks" to every atom in the sample simultaneously. It can deplete the ground state of the entire ensemble, leading to the power broadening we'll review in a moment.

But for an inhomogeneously broadened line, a monochromatic laser, by its very nature, is only resonant with a very specific, narrow subset of the atoms—the velocity class whose Doppler shift brings them into perfect resonance. The laser will interact strongly with this small group, but leave the vast majority of other atoms completely untouched.

Understanding this distinction is the absolute key that unlocks the concepts of "hole burning" and "Lamb dips." It is this selective interaction that we will exploit to defeat Doppler broadening. So, before we jump into that, let's quickly refresh our memory on how saturation affects a simple, homogeneous line.

Page 8:

Alright, let's do a quick review of power broadening for a homogeneously broadened transition. You'll recall we derived this in detail earlier in the course, in Section 3.6, but the key result is what's important here.

The central equation describes how the width of the transition changes with laser intensity. It reads:

Δ ω s = γ 1 + S 0

$$\Delta\omega_{\mathrm{s}} = \gamma\sqrt{1 + S_0}$$

Let's break this down term by term.

First, Delta omega sub s. This is the new, *saturated* half-width at half-maximum, or HWHM, of our spectral line. This is the width we measure in the presence of the intense laser.

Next is gamma, the lowercase Greek letter. This is the *unsaturated* homogeneous half-width at half-maximum. This is the line's "natural" width, determined by spontaneous decay and collisions, in the limit of very low laser intensity.

Finally, we have the crucial term, S naught. S sub zero, or S naught, is the on-resonance saturation parameter. This is a dimensionless quantity that tells us how intense our laser is compared to the intensity needed to significantly affect the atomic populations. A value of S naught much less than one means we're in the weak-field limit. A value of S naught of one or greater means we are in the saturation regime. We'll define S naught formally on the next couple of slides, but for now, just think of it as a knob that controls the laser intensity.

So, this equation tells us that as the laser intensity S naught increases from zero, the linewidth Delta omega sub s grows from its minimum value, gamma, becoming broader and broader.

Page 9:

So what is the physical meaning behind this mathematical formula for power broadening?

The first bullet point gets to the heart of it. An intense laser field drives transitions from the ground state to the excited state very, very quickly. If the laser is strong enough, it can pump atoms out of the ground state faster than they can relax back down. This leads to a significant depletion of the ground-state population.

From the perspective of the atom's interaction with the light field, this rapid cycling shortens the effective lifetime of the atom in a given state. And as we know from the Heisenberg uncertainty principle, a shorter lifetime in a state corresponds to a larger uncertainty in its energy. This energy uncertainty manifests directly as a broadening of the spectral line. So, power broadening is a direct consequence of the uncertainty principle applied to a strongly driven system.

Now, a key feature of this process for a *homogeneous* system is that the new, power-broadened lineshape is still a perfect Lorentzian. It's just a Lorentzian with an enlarged width, Δ ω s $\Delta\omega_s$. The fundamental character of the line doesn't change; it just gets wider and, as a consequence of conserving the total area under the curve in some contexts, it also gets shorter.

Page 10:

This Paage provides a wonderful visual illustration of the power broadening effect we've just discussed. We're looking at four plots, each showing the absorption as a function of frequency detuning from resonance. In each plot, the blue curve represents the low-intensity, unsaturated profile, while the red curve shows the high-intensity, saturated profile. The only thing changing between the four plots is the value of the on-resonance saturation parameter, S 0 S_0, which is controlled by the slider below each graph.

Let's start with the top-left plot, where the saturation parameter S 0 S_0 is set to zero. As we'd expect, the high-intensity red curve lies perfectly on top of the low-intensity blue curve. There is no saturation, and the half-width at half-maximum is simply γ γ.

Now, look at the top-right plot. Here, S 0 S_0 has been increased to 0.5 0.5. The red curve is now visibly different. Notice two things: its peak height is lower than the blue curve, and its width is broader. This is power broadening in action.

Moving to the bottom-left, S 0 S_0 is now 2.0 2.0. We are well into the saturation regime. The effect is even more dramatic. The red Lorentzian is significantly shorter and wider than the original blue profile.

Finally, in the bottom-right plot, we have an S 0 S_0 of 9.0 9.0, representing deep saturation. The red line is now very broad and flat. We have strongly depleted the absorbing power of the medium.

These plots perfectly visualize the formula from the previous slide, showing how the Lorentzian profile broadens and flattens as the laser intensity, quantified by S 0 S_0, is increased.

Page 11:

Now, let's formally define this crucial saturation parameter. We've been using S naught, the on-resonance value, but the saturation parameter, capital S, is actually a function of the laser frequency, omega. The definition given here is for a single-frequency, continuous-wave laser, which is exactly the tool we use for this kind of high-resolution spectroscopy.

The formal definition is given by the equation:

S (ω) = I (ω) I s a t · (γ 2) 2 (ω − ω 0) 2 + (γ 2) 2 .

$$S(\omega) = \frac{I(\omega)}{I_{\text{sat}}} \cdot \frac{\left(\frac{\gamma}{2}\right)^2}{(\omega - \omega_0)^2 + \left(\frac{\gamma}{2}\right)^2}.$$

Let's deconstruct this important expression.

The first term, I of omega, is the spectral power density of our laser at the angular frequency omega. You can think of this simply as the laser's intensity. Its units would be something like Watts per square meter.

The denominator of that first fraction contains I sat. This is the saturation intensity. This is an incredibly important parameter. It is a property of the *atomic transition itself*, not the laser. It represents the intensity required to cause significant saturation. We'll define it more physically on the next slide.

The second part of the expression is a normalized Lorentzian lineshape function. Omega is the variable laser frequency, and omega naught is the center resonance frequency of the atom.

This Lorentzian factor tells us that saturation is most effective when the laser is tuned directly to the atomic resonance, where omega equals omega naught. As we tune the laser away from resonance, the saturation effect falls off rapidly, following this Lorentzian shape.

Page 12:

Let's continue breaking down the components of our saturation parameter and see what it simplifies to on resonance.

First, let's revisit I sat I_{sat}, the saturation intensity. What does it physically mean? It is defined as the intensity at which the rate of stimulated transitions (absorption and stimulated emission) becomes equal to the total relaxation rate of the population difference. This relaxation rate includes spontaneous decay and any dephasing processes like collisions. So, you can think of I sat I_{sat} as the point of "fair competition": it's the laser intensity at which the laser's influence on the atom becomes just as strong as the atom's natural tendency to relax back to equilibrium.

Next, ω 0 ω_0, as we've said, is the center, natural transition frequency for a particle at rest.

The third bullet point simply reiterates a key concept: The Lorentzian line-shape factor in the full expression for S (ω) $S(\omega)$ ensures that the saturation effect is strongest when the laser frequency ω ω is tuned exactly to the atomic resonance ω 0 ω_0.

And this brings us to the final point. What happens on exact resonance, when ω = ω 0 $\omega = \omega_0$? In this case, the Lorentzian term in our definition becomes exactly one. The expression simplifies dramatically, and we recover the on-resonance saturation parameter, which we call S 0 S_0. So, S 0 S_0, which is what we used in our power broadening formula, is simply S (ω 0) $S(\omega_0)$, and it's given by the ratio of the on-resonance laser intensity, I 0 I_0, to the saturation intensity of the transition, I sat I_{sat}.

This S 0 S_0, the ratio I 0 / I sat I_0/I_{sat}, is the dimensionless parameter we will use most often to characterize the strength of our saturation.

Page 13:

Now that we have a firm grasp on the meaning of S 0 S_0, let's establish a practical rule of thumb for interpreting its value. This is extremely useful when you're in the lab setting up an experiment.

First, if S 0 S_0 is much, much less than 1 1… that is, the laser intensity is far below the saturation intensity… we are in the weak-field limit. In this regime, the laser acts as a small perturbation. The ground-state population is barely affected, so there is negligible depletion. The absorption of the medium is linear with the laser intensity. This is the regime of linear optics.

Next, when S 0 S_0 is approximately equal to 1 1. This signals the onset of saturation. At this point, the laser is strong enough to cause a significant redistribution of the atomic populations between the ground and excited states. The nonlinear effects we’re interested in, like power broadening, become clearly observable. This is the transition point into the nonlinear optics regime.

Finally, if S 0 S_0 is much, much greater than 1 1, we are in the deep saturation, or strong-field, limit. Here, the laser field is so intense that it overwhelmingly dominates the dynamics. The transition is effectively “bleached,” meaning the ground state is so depleted that the medium becomes nearly transparent to the resonant light. In this regime, power broadening is the dominant feature of the lineshape.

So, keep these three regimes in mind: S 0 S_0 much less than 1 1 is linear, S 0 S_0 around 1 1 is the interesting onset of nonlinearity, and S 0 S_0 much greater than 1 1 is the heavily saturated, power-broadened limit.

Page 14:

To set the stage for our main topic, let's quickly write down the familiar expressions for the absorption coefficient of a homogeneous medium, both with and without saturation. Think of this as our reference point.

First, let's consider unsaturated absorption. This is the case where the laser is weak, so S 0 S_0 is close to zero. The absorption coefficient, which we'll call α 0 (ω) $\alpha_0(\omega)$, is given by the expression:

α 0 (ω) = N σ 0 · (γ 2) 2 (ω − ω 0) 2 + (γ 2) 2 .

$$\alpha_0(\omega) = N\,\sigma_0 \cdot \frac{(\gamma/2)^2}{(\omega-\omega_0)^2+(\gamma/2)^2}.$$

Let's define the terms.

N N is the ground-state population density, for example, the number of atoms per cubic centimeter.

σ 0 σ_0 is the peak absorption cross-section at resonance. It tells you the effective area an atom presents to a photon right at the center of the transition.

The rest of the expression is simply the normalized Lorentzian lineshape. This entire formula is essentially a more detailed version of Beer's Law, accounting for the frequency dependence of the interaction.

Now, what happens when we introduce saturation?

Page 15:

When we turn up the laser intensity and introduce saturation, the expression for the absorption coefficient for our homogeneous medium takes on a beautifully simple form. The new, saturated absorption coefficient, α (ω) $\alpha(\omega)$, is given by:

α (ω) = α 0 (ω) 1 + S (ω) .

$$\alpha(\omega) = \frac{\alpha_0(\omega)}{1 + S(\omega)}.$$

Let's think about the physics here. The denominator, 1 + S (ω) $1 + S(\omega)$, captures the effect of "bleaching" the medium. As the laser intensity, and thus S (ω) $S(\omega)$, increases, the denominator gets larger, and the overall absorption α (ω) $\alpha(\omega)$ decreases. This makes perfect sense: as we pump more atoms into the excited state, there are fewer atoms remaining in the lower state available to absorb light. The medium becomes more transparent.

Now for a crucial insight. The unsaturated absorption, α 0 (ω) $\alpha_0(\omega)$, has a Lorentzian lineshape. And the saturation parameter, S (ω) $S(\omega)$, also has a Lorentzian frequency dependence. When you divide a Lorentzian by one plus another Lorentzian, the resulting mathematical form is *also* a Lorentzian. However, this new Lorentzian is both broader and shallower than the original one. This mathematical result is the precise origin of power broadening for a homogeneous line. It confirms exactly what we saw in the diagrams a few slides ago.

Okay. These expressions for the homogeneous case are our essential baseline. We now have all the tools we need to transition to the much more interesting, and more complex, Doppler-broadened, inhomogeneous situation.

Page 16:

Alright. Let’s take a moment. We have thoroughly reviewed the behavior of a homogeneously broadened line under saturation. We’ve seen that it simply gets wider and shorter, but remains a Lorentzian.

Now, we pivot to the central topic of this lecture. What happens when we apply this same intense, monochromatic laser beam to an *inhomogeneously* broadened system, like a Doppler-broadened gas? As we will see, the result is completely different, and far more interesting.

Page 17:

So, let's officially make the transition to the inhomogeneous case and remind ourselves why Doppler broadening is such a dominant effect in gases.

At room temperature, the thermal kinetic energy of atoms and molecules is significant. This translates into typical thermal velocities on the order of 100 to 1000 meters per second. These are not trivial speeds.

Now consider a molecule moving with a velocity component v z v_z along the axis of a laser beam, which we'll call the z-axis. Due to the Doppler effect, the frequency of the light that this molecule experiences in its own rest frame, which we'll call ω ′ ω', is shifted from the laser's frequency in the lab frame, ω ω. The relationship is given by:

ω ′ = ω − k v z .

$$\omega' = \omega - k\, v_{\mathrm{z}}.$$

Here, k k is the magnitude of the laser's wave-vector, which is equal to ω / c ω/c. And v z v_{z} is that velocity component. The sign convention here means that a molecule moving towards the laser source (with a negative v z v_{z}) sees an up-shifted frequency, while one moving away (positive v z v_{z}) sees a down-shifted frequency. This shift is the fundamental source of the broadening.

Page 18:

The velocity component along the laser axis, v z v_{z}, is not a single value; rather, the molecules in a gas have a distribution of velocities. For a gas in thermal equilibrium, this distribution follows the Maxwell-Boltzmann statistics. The probability of finding a molecule with a velocity component between v z v_{z} and v z + d v z $v_{\mathrm{z}} + dv_{\mathrm{z}}$ is given by the function:

f (v z) d v z = 1 v p π exp [− (v z v p) 2] d v z .

$$f(v_{\mathrm{z}})\, dv_{\mathrm{z}} = \frac{1}{v_{\mathrm{p}}\sqrt{\pi}} \exp\left[-\left(\frac{v_{\mathrm{z}}}{v_{\mathrm{p}}}\right)^{2}\right] dv_{\mathrm{z}}.$$

Here, the term v p v_{p} is the most probable speed, and it's given by the square root of two times the Boltzmann constant, k B k_{B}, times the temperature, T T, all divided by the mass of the molecule, m m.

v p = 2 k B T m .

$$v_{\mathrm{p}} = \sqrt{\frac{2\,k_{\mathrm{B}}\,T}{m}}.$$

This function, as you can see from its form, is a Gaussian.

Now, since the absorption of light depends on the Doppler-shifted frequency, and the Doppler shift depends on v z v_{z}, the overall absorption profile of the gas is found by integrating the contributions from all velocity classes. Because the velocity distribution is a Gaussian, the resultant absorption profile is also a Gaussian. The full width at half-maximum, or FWHM, of this Doppler-broadened profile is given by the famous formula for the Doppler width, which we denote as Δ ω D $\Delta\omega_{\mathrm{D}}$:

Δ ω D = ω 0 c 8 k B T ln 2 m .

$$\Delta\omega_{\mathrm{D}} = \frac{\omega_0}{c}\sqrt{\frac{8\,k_{\mathrm{B}}\,T\,\ln 2}{m}}.$$

This formula tells us that the Doppler width is larger for lighter atoms, higher temperatures, and higher transition frequencies.

Page 19:

Here is the crucial point that motivates our entire discussion of saturation spectroscopy. This Doppler width, Δ ω D $\Delta\omega_D$, which we just defined, often exceeds the homogeneous width, γ γ, by orders of magnitude. For a typical visible transition in an atomic vapor at room temperature, the Doppler width might be a gigahertz, that's 10 9 H z $10^9\,\mathrm{Hz}$, while the natural homogeneous width might only be ten megahertz, 10 7 H z $10^7\ \mathrm{Hz}$. That’s a factor of a hundred!

The staggering implication is that any interesting spectral features that are narrower than the Doppler width, like hyperfine splittings or the true natural linewidth itself, are completely hidden and washed out. They are buried under this enormous Doppler "blanket".

This leads to the conclusion in the slide: sub-natural features are completely inaccessible unless we employ special techniques. And the most powerful of these techniques is saturation spectroscopy, which is designed precisely to pull back this Doppler blanket and reveal the sharp features hiding underneath.

Page 20:

So, how do we begin to defeat this enormous Doppler broadening? The key insight lies in recognizing that a narrow-line laser does not interact with all the atoms in the gas equally.

Let's consider a monochromatic electromagnetic wave, represented by the electric field E = E 0 cos (ω t − k z) $E = E_0 \cos(\omega t - kz)$. This wave propagates through our gas of atoms. A given atom, which has a natural resonance frequency of ω 0 ω_0 and a homogeneous half-width of γ γ, will only interact strongly with this laser if the frequency it *sees* in its own moving frame is very close to its resonance frequency.

Mathematically, this leads to the condition shown on the slide:

| ω − ω 0 − k v z | ≤ γ

$$|\omega - \omega_0 - kv_z| \leq \gamma$$

Let's unpack this. The term ω − k v z $\omega - kv_z$ is the Doppler-shifted laser frequency as seen by the atom. For strong interaction, this perceived frequency must fall within the homogeneous linewidth, γ γ, of the atom's own resonance frequency, ω 0 ω_0.

In other words, a highly monochromatic laser selectively excites only those molecules whose velocity, v z v_z, shifts them into resonance. This creates what we can call an effective "selectivity window" in velocity space. The laser picks out just one small group of atoms from the entire thermal distribution.

Page 21:

This velocity-selective interaction defines a "selectivity window" in velocity space whose width we can now calculate.

From the resonance condition on the previous slide, we can see that the range of velocities that are efficiently excited is given by rearranging the inequality. This gives us a velocity width, which we'll call Δ v z Δv_z, that is approximately equal to γ / k γ/k.

So, the width of the velocity class we select is directly proportional to the homogeneous linewidth, γ γ, and inversely proportional to the wave vector of the light, k k. Since γ γ is typically very small (on the order of megahertz) and k k is large for visible light, this velocity window, Δ v z Δv_z, is extremely narrow compared to the overall thermal velocity distribution.

The crucial consequence is the final point on the slide: all other velocity classes, that is, all the atoms outside this narrow window, are far from resonance. The laser light is so far from their Doppler-shifted transition frequency that they hardly absorb at all. They are effectively invisible to the laser, and the laser is invisible to them.

Page 22:

This diagram provides a perfect visual summary of the concept of velocity-selective excitation.

What we are looking at is a plot where the horizontal axis represents the velocity component, v v, along the laser beam, and the vertical axis represents the number of molecules, N (v) $N(v)$.

The broad blue curve is the Doppler-broadened profile. It's a Gaussian, representing the Maxwell-Boltzmann distribution of velocities in the gas. This represents all the atoms in our sample.

Now, imagine we shine in a narrow-band laser with a frequency ω ω that is tuned slightly above the rest-frame atomic resonance, ω 0 ω_0. According to our resonance condition, v z = ω − ω 0 k $v_z = \frac{\omega-\omega_0}{k}$, this laser will be resonant with atoms that have a specific positive velocity.

This is shown by the shaded red area labeled "Excited Velocity Group." The laser doesn't interact with the whole distribution. It only interacts with, or "talks to," this very narrow slice of the velocity distribution. The center of this slice is at v = ω − ω 0 k $v = \frac{\omega-\omega_0}{k}$, and its width, Δ v Δv, is approximately γ k $\frac{\gamma}{k}$.

This process is the first step towards overcoming Doppler broadening: we've found a way to address a single, well-defined group of atoms.

Page 23:

This paage simply provides a detailed text caption for the figure we just analyzed. I will summarize its key points to reinforce the concepts.

Figure 1 illustrates the interaction of a monochromatic laser with a gas of molecules. First, the broad blue curve represents the Maxwell–Boltzmann distribution of the molecular velocity component, v z v_z, along the laser beam. This is the origin of the inhomogeneous Doppler broadening of the absorption profile.

Second, a narrow-band laser with frequency ω ω is shone into the gas. It selectively interacts *only* with molecules in a very narrow velocity range, Δ v z Δv_z.

Which molecules are these? They are the ones for whom the Doppler shift brings their natural resonance frequency, ω 0 ω_0, into coincidence with the laser frequency, ω ω. This is described by the resonance condition we’ve been discussing: | ω − ω 0 − k v z | < γ

$$|\omega - \omega_0 - k v_z| < \gamma$$

where γ γ is the homogeneous linewidth.

This “selectivity window,” shown as the shaded red area in the diagram, is centered at a velocity v z = ω − ω 0 k .

$$v_z = \frac{\omega - \omega_0}{k}.$$

Its width is approximately Δ v z = γ k .

$$\Delta v_z = \frac{\gamma}{k}.$$

Crucially, all molecules outside this specific velocity class are far from resonance and are therefore not excited. This selective power is what we are about to exploit.

Page 24:

Now that we have the qualitative picture of velocity selection, let's write down the mathematics more formally. Let's consider the absorption cross-section for a *specific* velocity class.

The first point here gives the cross-section for a transition from a lower state, which we'll call ket 1, to an upper state, ket 2, for a molecule that has a specific axial velocity, v z v_z. The cross-section, which we denote as sigma sub one-two of (ω , v z) (ω, v_z), is a function of both the laser frequency and the molecule's velocity. It's given by the expression:

σ 12 (ω , v z) = σ 0 · (γ 2) 2 (ω − ω 0 − k v z) 2 + (γ 2) 2 .

$$\sigma_{12}(\omega, v_z) = \sigma_0 \cdot \frac{(\gamma/2)^2}{(\omega - \omega_0 - k\, v_z)^2 + (\gamma/2)^2}.$$

Let's look closely at this equation. It's a standard Lorentzian profile. However, the resonance condition in the denominator is not just ω − ω 0 $\omega - \omega_0$. It's ω − ω 0 − k v z $\omega - \omega_0 - k\, v_z$. This means that for a molecule with velocity v z v_z, the entire Lorentzian profile is shifted in frequency space by the Doppler shift, k v z $k\, v_z$. For each velocity class, there is a different, shifted Lorentzian. This is the mathematical essence of inhomogeneous broadening.

Page 25:

Let's clarify the terms in that cross-section formula and its physical interpretation.

First, sigma naught. This is the peak absorption cross-section. It's the maximum value of the cross-section, which occurs when the laser is tuned to the Doppler-shifted center of the line for that specific velocity class. That is, when ω = ω 0 + k v z $\omega = \omega_0 + kv_z$.

Next, gamma. As before, gamma is the Half-Width at Half-Maximum (HWHM) of the homogeneous Lorentzian. Physically, it's the total decay rate of the coherence between the two levels. For a simple two-level system, it's the sum of the population decay rates of the upper and lower levels, γ = γ 1 + γ 2 $\gamma = \gamma_1 + \gamma_2$.

So, the resonant velocity class, this narrow slice of atoms that can interact with the laser, really does act as a “thin slice” of the full Maxwell–Boltzmann distribution. We are probing just one piece of the whole puzzle.

The final point is critical for what comes next. Under weak probe conditions, meaning the saturation parameter S S is essentially zero, this Lorentzian formula we’ve written down describes the complete absorption probability for that single velocity slice. But what happens when we turn up the intensity, when S S is no longer zero? That’s when we start to burn a hole.

Page 26:

Alright. So we have firmly established that a monochromatic laser interacts with only a single, narrow velocity class within a Doppler-broadened profile.

The next logical question is: what happens if we make that laser beam *intense*? What if the saturation parameter S is no longer negligible?

This leads us directly to the central concept of this lecture: the phenomenon of "hole burning."

Page 27:

Let's build a qualitative picture of what we mean by "hole burning," a phenomenon also known as a "Bennet hole," after William R. Bennett Jr., who first predicted it.

Imagine we now use an intense, narrow-band laser. We'll call this the "pump" laser. Because it's intense, it has a significant saturation parameter, S > 1 $S > 1$. Because it's narrow-band, it is still velocity-selective. The laser will drive the transition from the ground state to the excited state, but *only* for those atoms in the resonant velocity class. This intense pumping depletes the ground-state population, which we can call N 1 (v z) $N_1(v_z)$, specifically and only for those resonant velocities.

Simultaneously, as these atoms are removed from the ground state, they must appear in the excited state. So, the pump laser simultaneously populates the excited state, N 2 (v z) $N_2(v_z)$, for that very same velocity class.

The result is a striking modification of the population distributions. If we were to plot the ground state population, N 1 N_1, as a function of velocity, v z v_z, we would see the original Maxwell-Boltzmann distribution but with a narrow dip, or "hole," burned into it at the resonant velocity. Conversely, if we plotted the excited state population, N 2 N_2, we would see zero population everywhere *except* for a narrow peak at that same resonant velocity.

And what is the width of this hole and this peak? It's not the natural homogeneous width γ γ. It is the *power-broadened* homogeneous width, which we will call γ s γ_s, because the hole is being created by a strong, saturating laser.

Page 28:

This diagram provides a stunning visualization of saturation hole burning. It plots the population, N (v z) $N(v_z)$, as a function of the axial velocity, v z v_z.

There are two curves shown.

The blue curve, labeled N 1 (v z) $N_1(v_z)$, represents the ground state population. You can see it largely follows the broad, Gaussian shape of the original Maxwell-Boltzmann distribution. However, at a specific velocity, labeled v z ′ v_z', there is a sharp, narrow dip. This is the "Bennett Hole." It is a direct visualization of the depletion of ground-state atoms in the velocity class that is resonant with our intense pump laser.

The red curve, labeled N 2 (v z) $N_2(v_z)$, represents the excited state population. This curve is essentially zero everywhere, as you would expect for a gas in thermal equilibrium. But, at the very same velocity, v z ′ v_z', there is a sharp, narrow peak. This "Excited State Peak" consists of the very atoms that were removed from the ground state to create the Bennett hole.

The diagram beautifully annotates the key features. The width of the hole and the peak is the power-broadened homogeneous width, γ s γ_{S}. This is much, much narrower than the overall Doppler width of the distribution, labeled Δ v D Δv_{D}. This picture perfectly captures the essence of velocity-selective saturation.

Page 29:

Let's now move from the qualitative picture to a quantitative description of this population change. We'll start by considering the case of weak saturation, where the saturation parameter S is much less than 1.

Our starting point is a general solution that comes from analyzing the rate equations for a two-level system under the influence of a laser field. This solution tells us how the population of the ground state changes. The equation is:

N 1 0 − N 1 = Δ N 1 0 1 + S .

$$N_1^0 - N_1 = \frac{\Delta N_1^0}{1+S}.$$

Let's break this down. The left side, N 1 0 − N 1 $N_1^0 - N_1$, is the change, or depletion, of the ground state population. Δ N 1 0 ΔN_1^0, in the numerator on the right, is the initial, unsaturated population *difference* between the ground and excited states. And S is the saturation parameter, which depends on both laser intensity and frequency.

For the case we're considering, weak saturation where S ≪ 1 $S \ll 1$, we can use a Taylor expansion. The expression simplifies to: The change in population is approximately equal to Δ N 1 0 (1 − S) $\Delta N_1^0(1-S)$. Wait, that's not quite right. A Taylor expansion of 1 / (1 + S) $1/(1+S)$ for small S is 1 − S $1-S$. So the change in population, N 1 0 − N 1 $N_1^0 - N_1$, should be approximately Δ N 1 0 ⋅ S $\Delta N_1^0 \cdot S$. Let's re-examine the slide. Ah, the slide seems to be calculating the *saturated population difference*, Δ N ΔN. Let's assume the slide meant to write the saturated population difference

Δ N = Δ N 0 1 + S .

$$\Delta N = \frac{\Delta N^0}{1+S}.$$

In the weak field limit S ≪ 1 $S \ll 1$, the change in population difference Δ N 0 − Δ N $\Delta N^0 - \Delta N$ is then

Δ N 0 S 1 + S

$$\frac{\Delta N^0 S}{1+S}$$

which is approximately Δ N 0 ⋅ S $\Delta N^0 \cdot S$. Let's proceed with the expression given on the slide, but with the understanding that this is a simplified result from a rate equation model. I will note this is a simplification. The key is the dependence on S.

Let's define the terms clearly.

N i 0 N_i^0 is the equilibrium population of level 'i' in the absence of the laser. It follows the Maxwell-Boltzmann distribution. Δ N 1 0 ΔN_1^0 is the unsaturated population *difference*, N 1 0 − N 2 0 $N_1^0 - N_2^0$. For most systems at thermal equilibrium, N 2 0 N_2^0 is nearly zero, so this is approximately just the total ground state population. This population difference is what drives the absorption process.

Page 30:

Let's continue with our quantitative description of the population change.

The linear approximation we discussed, where the depletion is proportional to the saturation parameter S, is generally valid as long as the depletion remains small, say less than about 10 percent. This corresponds to the regime where S is significantly less than 1.

Now for the most crucial point on this slide. The change in population is *velocity-selective*. Why? Because the saturation parameter, S, is itself strongly dependent on velocity. Recall the definition: S is a function of the detuning from resonance, which is ω − ω 0 − k v z $\omega - \omega_0 - k\, v_z$. This means that S, and therefore the population depletion, will have a large value only for velocities v z v_z that are very close to the resonant velocity. For all other velocities, S will be nearly zero, and the population will be unchanged.

Page 31:

To make this velocity- selectivity perfectly explicit, let's write down the full form of the saturation parameter as a function of both laser frequency ω ω and molecular velocity v z v_z.

The expression is:

S (ω , v z) = S 0 × (γ / 2) 2 (ω − ω 0 − k v z) 2 + (γ / 2) 2 .

$$S(\omega, v_z) = S_0 \times \frac{(\gamma/2)^2}{(\omega - \omega_0 - k\, v_z)^2 + (\gamma/2)^2}.$$

The numerator is (γ / 2) 2 $(\gamma/2)^2$. The denominator is the quantity (ω − ω 0 − k v z) 2 + (γ / 2) 2 $(\omega - \omega_0 - k\, v_z)^2 + (\gamma/2)^2$.

Let's interpret this in the context of hole burning.

S 0 S_0 here is the on- resonance saturation parameter. It represents the maximum possible saturation, which occurs for those molecules whose velocity v z v_z *exactly* satisfies the Doppler resonance condition, making the denominator of the Lorentzian minimal.

Now, if we fix the laser frequency ω ω and think of this expression as a function of v z v_z, we see that the saturation parameter itself has a Lorentzian shape in *velocity space*. The Lorentzian is centered on the resonant velocity, v z = (ω − ω 0) / k $v_z = (\omega - \omega_0)/k$. This Lorentzian envelope is what selects the "burning" region. It mathematically defines the position and shape of the Bennett hole in the velocity distribution.

Page 32:

So, what is the width of this hole that we burn in the velocity distribution?

If we look at the Lorentzian shape of the saturation parameter S S as a function of frequency, its half-width at half-maximum is given by our familiar power-broadening formula:

γ s = γ 1 + S 0 .

$$\gamma_{\mathrm{s}} = \gamma\sqrt{1 + S_0}.$$

This γ s γ_{s}, the power-broadened width, is the characteristic width of the Bennett hole in the frequency domain. It tells us the range of frequencies over which the saturation is significant.

This quantity, γ s γ_{s}, is often referred to as the "power-broadened homogeneous width." This is a very descriptive name. It's the intrinsic homogeneous width of the atom, γ γ, but made larger by the influence of the intense laser field, quantified by S 0 S_0. To find the width of the hole in velocity space, we would simply use the relation

Δ v s = γ s k .

$$\Delta v_{\mathrm{s}} = \frac{\gamma_{\mathrm{s}}}{k}.$$

Page 33:

We are now ready to write down the detailed expressions for the velocity-dependent populations, N one and N two, after the intense pump laser has been applied. These expressions might look a bit complicated, but they are just the combination of all the pieces we have assembled. We are assuming weak saturation, S ≪ 1 $S \ll 1$.

The new ground state population, N 1 (ω , v z) $N_1(\omega, v_\mathrm{z})$, is the original population, N 1 0 (v z) $N_1^0(v_\mathrm{z})$, *minus* a depletion term. This depletion term is a Lorentzian function of velocity. The expression is:

N 1 (ω , v z) = N 1 0 (v z) − Δ N 0 S 0 (γ 2) 2 γ 1 τ [(ω − ω 0 − k v z) 2 + (γ s 2) 2] .

$$N_1(\omega, v_\mathrm{z}) = N_1^0(v_\mathrm{z}) - \frac{\Delta N^0\, S^0\, (\gamma/2)^2}{\gamma_1\, \tau\, [(\omega - \omega_0 - k v_\mathrm{z})^2 + (\gamma_\mathrm{s}/2)^2]}.$$

Similarly, the new excited state population, N 2 N_2, is the original population, N 2 0 N_2^0, *plus* a peak term, which is also a Lorentzian function of velocity. The expression is analogous, but with a γ 2 τ $\gamma_2\tau$ factor in the denominator:

N 2 (ω , v z) = N 2 0 (v z) + Δ N 0 S 0 (γ 2) 2 γ 2 τ [(ω − ω 0 − k v z) 2 + (γ s 2) 2] .

$$N_2(\omega, v_\mathrm{z}) = N_2^0(v_\mathrm{z}) + \frac{\Delta N^0\, S^0\, (\gamma/2)^2}{\gamma_2\, \tau\, [(\omega - \omega_0 - k v_\mathrm{z})^2 + (\gamma_\mathrm{s}/2)^2]}.$$

Let's define the new symbols here.

Gamma one is the population decay rate out of the lower level, ket 1. In many cases, for a stable ground state, this is essentially zero unless there are collisions.

Gamma two is the population decay rate out of the upper level, ket 2. This is primarily due to spontaneous emission.

The term tau, which we'll define on the next slide, is the population relaxation time. Notice the width of these Lorentzians is the power-broadened width, γ s γ_s.

Page 34:

Let's define the relaxation times, tau and capital T, that appeared in the full rate equation solutions.

First, tau, the lowercase Greek letter. This is the longitudinal relaxation time, also known as the T1 time. It characterizes the timescale over which the *populations* of the states return to thermal equilibrium after being perturbed. For a simple two-level system, it's related to the individual population decay rates by:

τ = 1 γ 1 + 1 γ 2 .

$$\tau = \frac{1}{\gamma_1} + \frac{1}{\gamma_2}.$$

This can also be written as γ γ 1 γ 2 $\frac{\gamma}{\gamma_1\gamma_2}$, where γ γ is the sum of the rates.

Next, capital T. This is the transverse relaxation time, or T2 time. This is a measure of the coherence lifetime. It characterizes how quickly the relative phase of the superposition of the two states decays. It is related to the decay rates by:

Capital T equals 1 / (γ 1 + γ 2) $1/(\gamma_1 + \gamma_2)$, which is simply 1 / γ $1/\gamma$, our homogeneous linewidth parameter.

The final note on this slide points out a subtle but important detail. If the decay rates from the two levels, γ 1 γ_1 and γ 2 γ_2, are unequal, this will produce an asymmetry between the depth of the Bennett hole in the ground state and the height of the corresponding peak in the excited state. This is because the depletion and repopulation dynamics are governed by these different rates.

Page 35:

Let's summarize the key properties of the Bennett hole we've been discussing: its center, its width, and its depth.

First, the velocity of the hole's center. The hole is burned at the velocity class, v z ∗ v_z^*, that is brought into resonance by the Doppler effect. This velocity is given by rearranging the resonance condition:

v z ∗ = ω − ω 0 k

$$v_z^* = \frac{\omega - \omega_0}{k}$$

This equation shows us that we have direct experimental control over which velocity class we saturate. By simply tuning our laser frequency, ω ω, we can move the Bennett hole to any position we choose within the Doppler profile.

Second, the half-width of the hole in the frequency domain. As we've established, the hole's width is not the natural width γ γ, but the power-broadened homogeneous width, γ s γ_s, given by:

γ s = γ 1 + S 0

$$\gamma_s = \gamma\sqrt{1 + S_0}$$

Third, we can characterize the depth of the hole. A useful measure is the normalized depth, which is the fractional decrease of the population *difference* at the center of the hole.

Page 36:

Let's look at the formula for this normalized depth of the Bennett hole.

The expression on the slide represents the fractional change in the population difference, Δ N ΔN, evaluated right at the center of the hole, v z ∗ v_z^*. It is given by the equation:

\[\frac{\Delta N_0(v_\text{z}^{*)} - \Delta N(v_\text{z}^{*)}}{\Delta N_0(v_\text{z}^{*)}} \;=\; \frac{S_0}{1 + S_0}\,.\]

This is a beautiful and intuitive formula. It tells us how effective our saturation is. Let's consider the implication spelled out on the slide.

Suppose we set our laser intensity such that the on-resonance saturation parameter, S 0 S_0, is exactly equal to 1. Plugging this into the formula gives us 1 / (1 + 1) $1/(1+1)$, which is one-half. This means that when the laser intensity equals the saturation intensity, we have removed exactly 50 percent of the initial population difference at the center of the hole. This provides a very clear, physical meaning for the condition S 0 = 1 $S_0 = 1$. It's the point where you've bleached half of the resonant absorbers.

Page 37:

Here we have another clean, clear visualization of a Bennett hole in the velocity distribution.

The title "Bennet Hole in Velocity Distribution" describes exactly what we're seeing. The plot shows the ground state population, N 1 (v z) $N_1(v_z)$, as a function of the velocity component, v z v_z.

The dashed blue line represents the original, unperturbed Maxwell-Boltzmann velocity distribution. It's a smooth Gaussian.

The solid red line shows the population distribution *after* it has been interrogated by a strong, monochromatic laser beam. You can clearly see that a narrow "hole" has been burned into the distribution at a specific velocity, which here is around v z = 0.8 $v_z = 0.8$ in some arbitrary units.

This graph provides an unambiguous visual representation of the concept. Saturation doesn't affect the entire sample; it performs a kind of microscopic surgery, selectively removing atoms from one very specific velocity class, while leaving their neighbors untouched.

Page 38:

So, we've successfully created a hole in the velocity distribution of our atoms. This is a change in the microscopic properties of the gas. The crucial question now is: how does this microscopic change translate into a macroscopic, observable signal? Specifically, how does it affect the overall absorption of a laser beam passing through the gas?

To answer this, we first consider the contribution to the total absorption from just one infinitesimal velocity slice, d v z dv_z. The differential absorption coefficient, which we write as d α (ω , v z) $d\alpha(\omega, v_\mathrm{z})$, is given by the product of the number of absorbers in that slice and their cross-section. The number of effective absorbers is the population *difference*, Δ N (v z) $\mathit{\Delta} N(v_\mathrm{z})$. The cross-section is σ 1 − 2 (ω , v z) $\sigma_{1-2}(\omega, v_\mathrm{z})$. So, we have:

d α (ω , v z) d v z = Δ N (v z) σ 1 − 2 (ω , v z) d v z .

$$d\alpha(\omega, v_\mathrm{z})\, dv_\mathrm{z} = \mathit{\Delta} N(v_\mathrm{z})\, \sigma_{1-2}(\omega, v_\mathrm{z})\, dv_\mathrm{z}.$$

Here, Δ N (v z) $\mathit{\Delta} N(v_\mathrm{z})$ is the velocity-dependent population difference, which is N 1 (v z) − N 2 (v z) $N_1(v_\mathrm{z}) - N_2(v_\mathrm{z})$. And this is the quantity that now contains the Bennett hole we've just burned.

Page 39:

To find the total, macroscopic absorption coefficient, α (ω) $\alpha(\omega)$, which is what a photodetector would actually measure, we must sum up the contributions from all the individual velocity slices. In calculus terms, we must integrate the differential absorption coefficient over the entire range of velocities.

So, the total absorption coefficient is given by the integral:

α (ω) = ∫ − ∞ + ∞ Δ N (v z) σ 12 (ω , v z) d v z .

$$\alpha(\omega) = \int_{-\infty}^{+\infty} \mathit{\Delta}\, N(v_\mathrm{z})\, \sigma_{12}(\omega, v_\mathrm{z})\, dv_\mathrm{z}.$$

This integral runs over the entire Maxwell–Boltzmann distribution. Now, let’s think about what’s inside this integral. The population difference, Δ N (

v z) $\Delta N(v_z)$, now contains the Lorentzian-shaped hole that we burned. The cross-section, σ 12 (ω , v z) $\sigma_{12}(\omega, v_z)$, also has a Lorentzian dependence on v z v_z.

So, when we perform this integral, we are effectively convoluting these Lorentzian features with the overall Gaussian envelope of the velocity distribution. This combination of a Gaussian and Lorentzians mathematically leads to what is known as a Voigt profile. However, as we are about to see, a crucial approximation will simplify this picture dramatically.

Page 40:

Let's begin the process of evaluating this integral. The first step is to insert the explicit mathematical forms for the terms in the integrand. We will use the simplified model where the saturation affects the overall population difference.

The absorption coefficient, α (ω) $\alpha(\omega)$, is given by: α (ω) = Δ N σ 0 v p π

$$\alpha(\omega) = \frac{\Delta N\, \sigma_0}{v_p \sqrt{\pi}}$$

This prefactor contains the total population difference and peak cross-section. This is then multiplied by the integral from − ∞ $-\infty$ to + ∞ $+\infty$ of the following:

In the numerator, we have e − (v z / v p) 2 d v z $e^{-(v_z/v_p)^2}\, dv_z$. This is the Gaussian part from the Maxwell–Boltzmann distribution.

In the denominator, we have (ω − ω 0 − k v z) 2 + (γ s / 2) 2 $(\omega - \omega_0 - k\,v_z)^2 + (\gamma_s/2)^2$. This is the Lorentzian lineshape for the absorption, where we've used the power-broadened width, γ s γ_s, because we are considering the absorption of the saturating beam itself.

So, our task is to evaluate the integral of a Gaussian function multiplied by a Lorentzian function. This is the definition of a Voigt function.

Page 41:

Now for Step 2, which is the key intellectual leap that allows us to solve this integral easily. We must recognize the dominant region of the integrand.

The integrand is a product of a very broad Gaussian and a very narrow Lorentzian. The Lorentzian part, with its denominator containing (γ s / 2) 2 $(\gamma_\mathrm{s}/2)^2$, is only significant when the detuning is small. That is, when the absolute value of (ω − ω 0 − k v z) $(\omega - \omega_0 - kv_\mathrm{z})$ is less than or approximately equal to the power-broadened width, γ s γ_s.

This condition corresponds to a very narrow velocity interval, Δ v z Δv_z, which is equal to γ s / k γ_s/k.

Now, we make a crucial comparison. In almost any realistic gas-phase experiment, the Doppler width, δ ω D $\delta\omega_\mathrm{D}$, is vastly larger than this power-broadened homogeneous width, γ s γ_s. This is especially true for weak saturation, where S 0 < 1 $S_0 < 1$. This means the Gaussian exponential factor varies extremely slowly over the tiny velocity interval where the Lorentzian is non-zero.

Because the Gaussian is essentially constant over the region where the rest of the integrand matters, we can pull it outside the integral! We evaluate the exponential factor at the center of the Lorentzian peak, which occurs at the resonant velocity, v z ∗ = (ω − ω 0) / k $v_z^* = (\omega - \omega_0)/k$. This approximation is extremely accurate and simplifies the problem immensely.

Page 42:

After we pull the slowly-varying Gaussian factor outside the integral, what remains inside is a much simpler integral. This is Step 3.

We are left with the integral ∫ − ∞ ∞ d v z (ω − ω 0 − k v z) 2 + (γ s 2) 2 $\int_{-\infty}^{\infty} \frac{dv_z}{(\omega - \omega_0 - kv_z)^2 + \left(\frac{\gamma_s}{2}\right)^2}$.

This is a standard definite integral. It is the integral of a pure Lorentzian function. The result of this integration is well-known and can be found in any table of integrals. The answer is: π π, divided by the product of k k and the half-width, (γ s 2) $\left(\frac{\gamma_s}{2}\right)$.

So, the entire integral collapses into this simple algebraic term. All we have to do now is substitute this result back into our expression, along with the Gaussian factor we pulled out earlier, and simplify everything to get our final answer.

Page 43:

After substituting and simplifying, we arrive at the final formula for the saturated absorption coefficient for a single laser beam interacting with a Doppler-broadened medium. And the result is remarkably, perhaps surprisingly, simple.

The saturated absorption coefficient, which we'll call α s (ω) $\alpha_\mathrm{s}(\omega)$, is given by the expression boxed on the slide:

α s (ω) = α 0 (ω) 1 + S 0

$$\alpha_\mathrm{s}(\omega) = \frac{\alpha_0(\omega)}{\sqrt{1 + S_0}}$$

Let's pause and appreciate this. After a multi-step integration and a key approximation, the complex interaction boils down to this: the saturated absorption profile has the *exact same shape* as the original, unsaturated Gaussian profile. The only effect of the saturation is to reduce its overall amplitude by a constant scaling factor, one over the square root of 1 + S 0 $1 + S_0$.

The second equation on the slide just makes this explicit, showing that α 0 (ω) $\alpha_0(\omega)$ is proportional to the Gaussian exponential factor:

α 0 (ω) ∝ exp [− (ω − ω 0) 2 (0.6 δ ω D) 2]

$$\alpha_0(\omega) \propto \exp\left[-\frac{(\omega - \omega_0)^2}{(0.6\,\delta\omega_\mathrm{D})^2}\right]$$

The denominator is just a conversion factor between FWHM and the standard deviation of the Gaussian. The key takeaway is that the shape remains perfectly Gaussian.

Page 44:

This leads us to a critical observation, and a point that often confuses students initially.

Despite the fact that we have burned a very real, narrow hole in the *velocity distribution* of the atoms, **no spectral "hole" appears** in the macroscopic absorption profile. When we scan our single laser's frequency and measure the total absorption, we do not see a narrow dip. All we observe is an overall vertical scaling—a uniform reduction—of the Doppler profile.

What is the reason for this? The detector in our experiment averages over the contributions from all velocities. At any given laser frequency ω ω, we are indeed burning a hole in a specific velocity class. However, this hole affects only a tiny fraction of the total number of molecules in the gas. As we scan the laser to a new frequency, we simply burn a new hole in a *different* velocity class. The local depletion is always there, but its contribution to the *total* attenuation is minuscule. When integrated over all velocities, the net effect is simply a slight reduction in the overall absorption, while the profile retains its Gaussian shape.

Page 45: Saturated Absorption of a Doppler-Broadened Line

This pair of interactive plots perfectly illustrates the conclusion we just reached. The title is "Saturated Absorption of a Doppler-Broadened Line." We are plotting the absorption coefficient versus frequency detuning. The solid blue line is the unsaturated profile, and the dashed red line is the saturated profile.

On the left, the saturation parameter S 0 S_0 is set to 0 0. As expected, the saturated and unsaturated profiles are identical. There is no effect.

Now, look at the plot on the right, where the saturation parameter S 0 S_0 has been cranked up to 3.0 3.0. Look closely at the red dashed line. It is still a perfect Gaussian. It has the same center and the same width as the blue unsaturated profile. The only difference is that its peak amplitude is significantly lower. The absorption has been uniformly suppressed across the entire profile, exactly as our formula α s = α 0 1 + S 0 $\alpha_s = \frac{\alpha^0}{\sqrt{1+S_0}}$ predicts.

The key observation, summarized in the box below, is the crucial takeaway: A single traveling laser wave interacting with an inhomogeneously broadened medium does *not* create a spectral hole or a Lamb Dip in the macroscopic absorption profile. The detector averages over everything. While a narrow group of atoms is saturated, their contribution is small compared to the whole. The result is just a uniform reduction of the entire Gaussian profile.

Page 46:

Let's just reiterate exactly why this Bennett hole is invisible in a simple, single-laser scan. There are three key reasons that work together.

First, the saturation hole itself occupies a frequency width, γ s γ_{s}, that is much, much smaller than the overall Doppler width. The hole is a tiny feature on a massive background.

Second, the process of measuring the macroscopic absorption involves an integration over all velocities. This integration process effectively “smears out” the effect of the local depletion over the full Gaussian envelope. At each laser frequency, you’re probing a different velocity class, so you never get to see the persistent effect of a hole at one specific spectral location.

Third, and most formally, the scaling factor we derived, which is (1 + S 0) − 1 / 2 $(1 + S_0)^{-1/2}$, is constant with respect to the laser frequency, ω ω, assuming our laser intensity is stable as we scan. A detector measuring the transmitted light cannot distinguish between a reduction in absorption due to saturation and, for example, simply having a lower concentration of atoms in the cell. It just sees a weaker signal across the board.

So, how can we ever hope to see this hole?

Page 47:

The solution, and the key to all of modern saturation spectroscopy, is elegant and powerful.

Instead of using one laser beam to both saturate and measure, we introduce a *separate probe laser* to interrogate the burned velocity class selectively. This is the foundation of what we call pump-probe spectroscopy.

The idea is simple: We use one laser, the strong "pump," to do the work of burning the hole in a specific, fixed velocity class. Then we use a second, weak "probe" laser, whose frequency we scan, to act as a reporter. The probe's job is simply to measure the absorption profile of the medium that has been prepared by the pump. When the probe laser's frequency is tuned to be resonant with the same velocity class that the pump has already depleted, it will experience reduced absorption. It will "see" the hole.

Page 48:

Let's outline this two-laser pump-probe configuration to detect the Bennett hole.

First, we have the pump laser. It has a wave-vector k 1 k_1 and a frequency ω 1 ω_1. This laser is strong, meaning it has a significant saturation parameter. Crucially, its frequency is held *fixed* at some value, usually somewhere within the Doppler profile.

The job of this fixed-frequency pump laser is to burn a Bennett hole in one specific velocity class. The velocity class that gets saturated, v z ∗ v_z^*, is the one that satisfies the resonance condition. Following the convention on the slide, this is given by:

v z ∗ =ω 0 −ω 1 k 1 .

$$v_z^* = \frac{\omega_0 - \omega_1}{k_1}.$$

So, by choosing the pump frequency ω 1 ω_1, we select and burn a hole in a single, well-defined velocity class.

Next, we have the probe laser. It has a wave-vector k 2 k_2 and a frequency ω ω. This laser must have a weak intensity, with a saturation parameter S ≪ 1 $S \ll 1$, to ensure that it doesn't cause any significant saturation itself. Its purpose is only to measure, not to perturb. The frequency, ω ω, of this probe laser is then scanned across the absorption profile to measure any changes in absorption caused by the pump.

Page 49:

So, under what condition will our probe laser "see" the hole burned by the pump?

The probe experiences reduced absorption when it is tuned to interrogate the very same velocity class, v z ∗ v_z^*, that has been saturated by the pump.

The pump selected the class v z ∗ = (ω 0 − ω 1) / k 1 $v_z^* = (\omega_0 - \omega_1)/k_1$. The probe, as we scan its frequency ω ω, will be resonant with a velocity class v = (ω 0 − ω) / k 2 $v = (\omega_0 - \omega)/k_2$.

The probe will see the dip when these two velocities are the same. By setting them equal and rearranging, we arrive at the condition for the spectral dip:

ω = ω 0 ± (ω 1 − ω 0) (k 1 k 2) .

$$\omega = \omega_0 \pm (\omega_1 - \omega_0)\left(\frac{k_1}{k_2}\right).$$

The plus or minus sign here is important and depends on the geometry of the experiment, specifically whether the pump and probe beams are co-propagating or counter-propagating. We'll clarify this on the next slide. This equation tells us precisely where in the spectrum the Bennett hole will appear.

Page 50:

Let's clarify the meaning of the ± $\pm$ sign in our condition for the spectral dip. The choice of sign depends on the relative directions of the pump and probe beams.

You use the + $+$ sign for co-propagating beams. This is the case where the pump laser and probe laser are traveling in the same direction, so their wave-vectors, k 1 k_1 and k 2 k_2, have the same sign.

You use the − $-$ sign for counter-propagating beams. This is where the pump and probe travel in opposite directions through the sample, so their wave-vectors have opposite signs, for instance, k 1 k_1 is positive and k 2 k_2 is negative.

This distinction arises because a given velocity class, say atoms moving with positive v z v_z, will have its resonance Doppler-shifted in one direction for a laser beam coming from the left, and in the opposite direction for a laser beam coming from the right. To have both beams interact with the same atoms, their frequencies must be tuned accordingly.

Page 51:

This diagram provides an excellent schematic of a pump-probe spectroscopy experiment designed to detect a Bennett hole.

We see an atomic vapor cell in the center. An intense pump laser, shown in red with frequency ω 1 ω_1 and wave-vector k 1 k_1, passes through the cell from left to right. Its frequency is fixed. This laser's job is to create the Bennett hole. As noted in box 1, it selectively interacts with and saturates atoms whose velocity vector v v satisfies the resonance condition, k 1 · v = ω 1 − ω 0 $k_1 \cdot v = \omega_1 - \omega_0$.

A second, weak probe laser, shown in blue with frequency ω ω and wave-vector k 2 k_2, is also passed through the cell and its transmission is monitored by a detector. The key is that the probe laser's frequency, ω ω, is scanned.

As described in box 2, the probe laser's absorption depends on atoms satisfying its own resonance condition, k 2 · v = ω − ω 0 $k_2 \cdot v = \omega - \omega_0$. When the scanned frequency ω ω reaches the value where the probe interacts with the *same velocity class* that has already been saturated by the pump, there are fewer ground state atoms available to absorb the probe light. The detector therefore measures a sharp dip in absorption. This dip *is* the spectral signature of the Bennett hole.

Page 52:

Now we arrive at the full integral expression for the probe absorption in the presence of the pump. This formula may look daunting, but it is a complete mathematical description of the pump-probe experiment. It represents the total absorption coefficient, alpha sub s, which is a function of both the fixed pump frequency, omega 1, and the scanned probe frequency, omega.

The expression is:

α s = σ 0 Δ N 0 v p π ∫ − ∞ ∞ exp ⁡ (− (v z / v p) 2) [1 − S 0 (γ / 2) 2 (ω 0 − ω 1 − k 1 v z) 2 + (γ s / 2) 2] 1 (ω 0 − ω − k 2 v z) 2 + (γ / 2) 2 d v z .

$$\alpha_{\mathrm{s}} = \frac{\sigma_0\,\Delta N_0}{v_{\mathrm{p}}\sqrt{\pi}}\int_{-\infty}^{\infty}\exp\left(-\left(v_{\mathrm{z}}/v_{\mathrm{p}}\right)^2\right)\left[1 - \frac{S_0\,(\gamma/2)^2}{(\omega_0-\omega_1-k_1 v_{\mathrm{z}})^2+(\gamma_{\mathrm{s}}/2)^2}\right]\frac{1}{(\omega_0-\omega-k_2 v_{\mathrm{z}})^2+(\gamma/2)^2}\,dv_{\mathrm{z}}.$$

The integrand itself has three parts multiplied together.

First, exp ⁡ (− (v z / v p) 2) $\exp\left(-\left(v_{\mathrm{z}}/v_{\mathrm{p}}\right)^2\right)$, which is our Gaussian velocity distribution.

Second, a population factor, which is [1 − S 0 (γ / 2) 2 (ω 0 − ω 1 − k 1 v z) 2 + (γ s / 2) 2] $\left[1 - \frac{S_0\,(\gamma/2)^2}{(\omega_0-\omega_1-k_1 v_{\mathrm{z}})^2+(\gamma_{\mathrm{s}}/2)^2}\right]$. This term is a Lorentzian describing the hole burned by the pump. It looks like 1 − S 0 (γ / 2) 2 (ω 0 − ω 1 − k 1 v z) 2 + (γ s / 2) 2 $1 - \frac{S_0\,(\gamma/2)^2}{(\omega_0-\omega_1-k_1 v_{\mathrm{z}})^2+(\gamma_{\mathrm{s}}/2)^2}$. This (1 − …) $(1 - \cdots)$ factor represents the fraction of the population *remaining* after the pump has acted.

Third, a probe interaction factor, which is another Lorentzian. This is 1 (ω 0 − ω − k 2 v z) 2 + (γ / 2) 2 $\frac{1}{(\omega_0-\omega-k_2 v_{\mathrm{z}})^2+(\gamma/2)^2}$. This describes the probability of the weak probe interacting with each velocity class.

So, the integral sums up, for every velocity slice, the product of: how many atoms are there (Gaussian), how many of them are left after the pump hits them (the hole factor), and how strongly the probe interacts with them (the probe Lorentzian).

Page 53:

Let's break down the physical meaning of the different parts of that complex integral from the previous page.

The first key piece is a Lorentzian factor that describes the interaction of the weak probe with the atoms. This Lorentzian is a function of the probe frequency ω ω and the velocity v z v_z. As we scan ω ω, this Lorentzian effectively sweeps through the velocity distribution, selecting which velocity class is being probed at any given moment.

The second key piece is the bracketed term, the [1 − …] $[1 - \cdots]$ factor. This term represents the population reduction produced by the strong pump laser. It's equal to 1 for velocities far from the pump's resonance, and it dips to a minimum at the velocity that is resonant with the pump. This term mathematically represents the Bennett hole.

The crucial part of this term is the *second* Lorentzian, the one inside the bracket. It ensures that the population depletion only occurs for velocities that satisfy the pump's resonance condition. The entire integral will only yield a significant signal change when the probe's Lorentzian and the pump's Lorentzian (the hole) overlap in velocity space.

Page 54:

After performing the velocity integration, which involves the same kind of approximation we made earlier—assuming the Lorentzians are much narrower than the Gaussian—we arrive at the final result for the pump-probe absorption spectrum.

The saturated absorption coefficient, α s (ω 1 , ω) $\alpha_{\mathrm{s}}(\omega_1, \omega)$, is given by the unsaturated Doppler profile, α 0 (ω) $\alpha_0(\omega)$, multiplied by a factor:

α s (ω 1 , ω) = α 0 (ω) [1 − Lorentzian dip term] .

$$\alpha_{\mathrm{s}}(\omega_1, \omega) = \alpha_0(\omega)[1 - \text{Lorentzian dip term}].$$

This Lorentzian term has a numerator proportional to S 0 S_0, and a denominator of the form (ω − ω ′) 2 + (Γ s / 2) 2 $(\omega - \omega')^2 + (\Gamma_{\mathrm{s}}/2)^2$.

So, the result is no longer a uniformly suppressed Gaussian. Instead, we see the original Doppler profile with a sharp, narrow Lorentzian dip subtracted from it. We have finally revealed the hole!

The central position of this dip, ω ′ ω', is given by the formula we derived earlier:

ω ′ = ω 0 ± (ω 1 − ω 0) k 1 k 2 .

$$\omega' = \omega_0 \pm (\omega_1 - \omega_0)\frac{k_1}{k_2}.$$

This tells us exactly where to look for the dip in our spectrum.

And what is the half-width of this observable dip? We'll define that on the next slide.

Page 55:

Let's now characterize the width and depth of this spectral dip that we observe in a pump-probe experiment.

The half-width of the dip, which we will call capital Gamma sub s, is given by:

Γ s = γ + γ s

$$\mathit{\Gamma}_{\mathrm{s}} = \gamma + \gamma_{\mathrm{s}}$$

Let's unpack this. Gamma is the unsaturated homogeneous width, which is the width of the Lorentzian describing the weak probe's interaction. Gamma sub s is the power-broadened homogeneous width, which is the width of the Bennett hole burned by the pump. The total width of the observed dip is the sum of these two widths, which arises from the convolution of the two Lorentzian profiles—the "hole" profile and the "probe" profile. We can write this explicitly as: gamma times the quantity, one plus the square root of (one plus S naught).

Γ s = γ (1 + 1 + S 0)

$$\mathit{\Gamma}_{\mathrm{s}} = \gamma\left(1 + \sqrt{1 + S_0}\right)$$

Now, what about the depth of the dip at its center, omega prime? In the limit of weak saturation, where S naught is much less than one, the change in absorption, Delta alpha, is given by:

Delta alpha at omega prime equals the unsaturated absorption at that frequency, alpha naught of omega prime, times S naught over four.

Δ α (ω ′) = α 0 (ω ′) S 0 4

$$\Delta\alpha(\omega') = \alpha_0(\omega')\frac{S_0}{4}$$

The depth of the dip is directly proportional to the saturation parameter, S naught.

And the key point, to summarize everything: The Bennett hole, which was invisible in a single-beam experiment, becomes observable as a sharp spectral dip because the probe laser allows us to specifically interrogate the depleted velocity group that was selected by the pump.

Page 56:

Now we move to a particularly clever and practical implementation of saturation spectroscopy known as the Lamb dip, named after Willis Lamb. This technique uses only a *single laser* but achieves the pump-probe effect by creating counter-propagating beams from that single laser.

The most common practical method is to take your laser beam, pass it through your sample cell, and then reflect it directly back upon itself with a mirror. This creates a standing wave inside the cell.

In this configuration, we have a forward-propagating wave with wave-vector k 1 k_1, and a backward-propagating wave with wave-vector k 2 k_2. Since it's a reflection, we have the condition: k 1 = − k 2

$$k_1 = -k_2$$

We can just call the magnitude k k.

Now, consider what different atoms experience. An atom moving with some velocity v z v_z will be Doppler-shifted into resonance with the forward wave at one laser frequency, and with the backward wave at a different frequency. But what about the special case of molecules that are not moving along the laser axis, that is, v z ≈ 0 $v_z \approx 0$?

These on-axis molecules experience *both* waves at essentially the same frequency, ω ω. They see the forward wave and the backward wave simultaneously.

This means the intensity experienced by this zero-velocity class of molecules is effectively doubled. This leads to a deeper saturation effect specifically for the atoms with v z ≈ 0 $v_z \approx 0$.

Page 57:

The unique saturation of the zero-velocity atoms has a profound effect on the measured absorption spectrum.

When we scan the frequency of our single laser, the resulting absorption profile shows a narrow dip precisely at the center of the transition, when ω = ω 0 $\omega = \omega_0$. This occurs despite the fact that we are still looking at the overall, broad Doppler envelope.

Let's think about why. When the laser is tuned far from the line center, say to the blue side, the forward-propagating wave interacts with atoms moving away from the laser, and the backward-propagating wave interacts with atoms moving towards the laser. These are two completely different velocity classes. Each class sees only one of the waves, so the saturation is moderate.

But when the laser is tuned to the exact line center, ω = ω 0 $\omega = \omega_0$, both the forward and backward waves become resonant with the *same* velocity class: the v z = 0 $v_z = 0$ class. This class sees both beams, experiences double the intensity, becomes more strongly saturated, and therefore absorbs less light. This reduction in absorption at the line center creates the "Lamb dip."

Page 58: This diagram provides a beautiful illustration of how a Lamb dip is formed by saturation in a standing wave.

We have a gas cell where an incident laser with wave-vector k k enters from the left, reflects off a mirror on the right, and propagates back as a reflected laser with wave-vector − k $-k$. The superposition of these two waves creates the green standing wave pattern.

Now let's look at the different atoms, represented by colored circles. Consider an atom moving to the left, with v z > 0 $v_z > 0$. As the label indicates, due to its Doppler shift, it can become resonant with the backward-propagating wave (the blue dashed line). Conversely, an atom moving to the right, with v z < 0 $v_z < 0$, can be resonant with the forward-propagating wave (the red dashed line). In both these cases, the atom interacts with only *one* of the traveling waves.

But now look at the yellow circle, the special molecule with v z ≈ 0 $v_z \approx 0$. This molecule is stationary with respect to the standing wave. It interacts with *both* the forward and backward waves. If it's located at an antinode of the standing wave, it experiences a very high intensity. As the label says, this leads to strong saturation. It is this enhanced saturation of the zero-velocity class, and only the zero-velocity class, that creates the Lamb dip.

Page 59:

Let's write down the mathematical expression for the population difference, Δ N $\mathit{\Delta N}$, in this standing-wave field.

The total population difference as a function of velocity, Δ N (v z) $\mathit{\Delta N}(v_z)$, is the initial population difference, Δ N 0 (v z) $\mathit{\Delta N}_0(v_z)$, multiplied by a saturation factor. This factor is now [1 − one hole term − a second hole term] $[1 - \text{one hole term} - \text{a second hole term}]$.

So, the expression is:

Δ N (v z) = Δ N 0 (v z) [1 − a first Lorentzian hole − a second Lorentzian hole] .

$$\mathit{\Delta N}(v_z) = \mathit{\Delta N}_0(v_z)\,[1 - \text{a first Lorentzian hole} - \text{a second Lorentzian hole}].$$

The first hole term is S 0 (γ / 2) 2 (ω 0 − ω − k v z) 2 + (γ s / 2) 2 $S_0 \frac{(\gamma/2)^2}{(\omega_0-\omega-kv_z)^2+(\gamma_s/2)^2}$. This is the hole burned by the forward-propagating wave, centered at a velocity v z = ω − ω 0 k $v_z = \frac{\omega-\omega_0}{k}$.

The second hole term is S 0 (γ / 2) 2 (ω 0 − ω + k v z) 2 + (γ s / 2) 2 $S_0 \frac{(\gamma/2)^2}{(\omega_0-\omega+kv_\mathrm{z})^2+(\gamma_\mathrm{s}/2)^2}$. This is the hole burned by the backward-propagating wave, which has a wave vector of minus k k. This hole is centered at a velocity v z = − ω − ω 0 k $v_\mathrm{z} = -\frac{\omega-\omega_0}{k}$.

So, for any off-resonant laser frequency ω ω, the standing wave burns *two* symmetric holes in the velocity distribution.

Page 60:

Now, what happens to these two holes as we tune our laser frequency? The two holes are centered at v z = + ω − ω 0 k $v_\mathrm{z} = +\frac{\omega-\omega_0}{k}$ and v z = − ω − ω 0 k $v_\mathrm{z} = -\frac{\omega-\omega_0}{k}$.

When the laser frequency ω ω is tuned to the exact atomic resonance, ω 0 ω_0, the term (ω − ω 0) $(\omega-\omega_0)$ becomes zero. This means both holes become centered at v z = 0 $v_\mathrm{z} = 0$. They perfectly overlap.

This overlap results in the deepest possible depletion, as the zero-velocity molecules are saturated by both the forward and backward waves simultaneously. This enhanced depletion of the zero-velocity molecules is precisely what leads directly to the formation of the Lamb dip in the absorption spectrum. The reduced absorption of this specific velocity class carves a narrow dip right at the center of the broad Doppler profile.

Page 61:

Let's look at the resulting absorption formula for the Lamb dip, specifically in the weak-field limit where S 0 S_0 is small.

The saturated absorption coefficient, α s (ω) $\alpha_\text{s}(\omega)$, is given by the unsaturated Doppler profile, α 0 (ω) $\alpha_0(\omega)$, multiplied by a dip factor. The expression is:

α s (ω) = α 0 (ω) [1 − S 0 2 1 1 + 4 (ω − ω 0) 2 γ s 2] .

$$\alpha_\text{s}(\omega) = \alpha_0(\omega)\left[1 - \frac{S_0}{2}\,\frac{1}{1 + \frac{4(\omega - \omega_0)^2}{\gamma_\text{s}^2}}\right].$$

Let's break this down. The first term, α 0 (ω) $\alpha_0(\omega)$, is just the peak Doppler absorption. This is multiplied by the dip-factor in the brackets. This factor is 1 minus a term that is sharply peaked when ω = ω 0 $\omega = \omega_0$. This "subtraction" is what creates the dip. The Lorentzian shape of this term means the dip itself will have a Lorentzian profile.

Now, let's look at the situation at the exact center of the resonance, when ω = ω 0 $\omega = \omega_0$. The Lorentzian term in the bracket becomes 1. The formula simplifies to:

α s (ω 0) = α 0 (ω 0) (1 − S 0) .

$$\alpha_\text{s}(\omega_0) = \alpha_0(\omega_0)\,(1 - S_0).$$

(Note: The formula on the slide uses 1 - S $_0$

. More detailed derivations often yield 1 - S $_0$ 2

if S$_0$ is for the total intensity. Assuming the slide's convention, the depth is directly proportional to S$_0$.)

This shows the absorption is reduced right at the line center.

Page 62:

Let's characterize the width and depth of this Lamb dip.

First, the half-width of the dip. The half-width at half-maximum of the Lamb dip, which we can denote as δ ω L D $\delta\omega_{LD}$, is simply equal to our old friend, the power-broadened homogeneous width, γ s γ_s. So,

δ ω L D = γ 1 + S 0 .

$$\delta\omega_{LD} = \gamma\sqrt{1 + S_0}.$$

This is a critical result. The width of the Lamb dip is determined by the homogeneous linewidth, *not* the Doppler width. This is why it's a sub-Doppler technique. By measuring the width of the Lamb dip, we can directly measure the homogeneous linewidth of a transition, even when it's buried inside a massive Doppler profile.

Now, what about the depth? As we saw, the depth is proportional to the saturation parameter, S 0 S_0. This presents a classic experimental trade-off. To get a deep, easily visible dip, you want to increase S 0 S_0 by turning up your laser power. However, as you increase S 0 S_0, you also increase the power broadening, which makes the dip wider.

A good practical compromise, as noted on the slide, is to choose a laser intensity such that S 0 S_0 is approximately equal to 1. This gives a significant reduction in absorption at the center—on the order of 50 percent—while avoiding the excessive power broadening that would occur at very high S 0 S_0.

Page 63:

Let's summarize the physical interpretation of the Lamb dip one more time to ensure the concept is crystal clear. We need to consider two distinct cases.

Case 1: For off-resonant frequencies, where the laser frequency ω ω is not equal to the atomic resonance ω 0 ω_0.

In this situation, the forward and backward traveling waves that make up the standing wave interact with two *different* and distinct velocity classes. The forward wave interacts with atoms at velocity v z = ω − ω 0 k $v_z = \frac{\omega-\omega_0}{k}$, while the backward wave interacts with atoms at velocity v z = − ω − ω 0 k $v_z = -\frac{\omega-\omega_0}{k}$.

Each of these velocity classes sees only a single traveling wave, and thus only experiences the intensity of that one wave. This leads to a moderate level of saturation for two separate groups of atoms.

Case 2: For exact resonance.

Page 64:

Now, let's consider the second case: what happens at exact resonance, when the laser frequency ω ω equals the atomic resonance ω 0 ω_0.

In this special case, the condition for resonance for both the forward and backward waves collapses to v z = 0 $v_z = 0$. This means that the *same* group of molecules—the zero-velocity class—interacts with *both* waves simultaneously.

For these molecules, the effective intensity they experience is the sum of the intensities of the two waves, which is essentially double the intensity seen by any off-resonant velocity class. This doubling of intensity leads to a much stronger saturation effect, which in turn causes a significant decrease in the absorption for this zero-velocity group.

Therefore, a sharp dip in the total absorption appears precisely at the line center. This feature is incredibly useful for practical applications like laser frequency locking with high precision, as the bottom of the dip provides a sharp, unambiguous marker for the true, un-shifted atomic resonance frequency.

Page 65:

This final diagram provides a comprehensive visual summary of Lamb Dip Spectroscopy.

The main graph at the top shows the macroscopic Absorption Profile, α (ω) $\alpha(\omega)$, as a function of frequency detuning. We see the broad, dashed-line Gaussian of the Doppler profile. Superimposed on this is the actual measured absorption, the solid blue line, which clearly shows the sharp, sub-Doppler Lamb Dip right at the center, ω = ω 0 $\omega = \omega_0$.

The two smaller diagrams at the bottom explain the origin of this dip by showing what’s happening in velocity space.

On the left, we see the "Off-Resonance" case. The forward and backward waves burn two separate, symmetric holes in the velocity distribution, well away from v z = 0 $v_z = 0$.

On the right, we see the "On-Resonance" case. Here, the two holes have coalesced. Both waves interact with the same v z = 0 $v_z = 0$ atoms, burning a single, deeper hole right at the center of the distribution. It is this enhanced saturation at v z = 0 $v_z = 0$ that is directly responsible for the Lamb dip we see in the frequency spectrum above. This is a perfect summary of the entire mechanism.

Page 66:

To solidify our understanding, let's make a direct comparison between saturating a homogeneous line versus an inhomogeneous line with a standing wave. The outcomes are qualitatively different.

First, consider homogeneous broadening, as depicted in Figure 2.7a which we will see on a later slide. In this case, a standing wave still creates a spatial modulation of the intensity, and therefore a spatial modulation of the saturation. This is called spatial hole burning. However, *every* molecule in the sample has the same resonance frequency. There is no velocity selection. The spectrum you observe is simply a power-broadened Lorentzian. The standing wave does not produce a deep, narrow dip in the frequency spectrum.

Now, consider inhomogeneous broadening, as depicted in Figure 2.7b. Here, velocity selection is the name of the game. As we've just seen, the special role of the zero-velocity class, which is the only class that can interact with both counter-propagating beams simultaneously, leads directly to the formation of a deep, narrow Lamb dip precisely at the un-shifted resonance frequency, ω 0 ω_0. It is the inhomogeneous nature of the broadening that makes this powerful technique possible.

Page 67

The slide also mentions an additional scenario for comparison, which corresponds to Figure 2.7c.

This scenario is our standard pump-probe experiment, but with a specific tuning. Here, the pump laser is fixed precisely at the line center, ω 0 ω_0. A separate, weak probe laser is then scanned across the profile.

In this case, the pump burns a Bennett hole in the v z = 0 $v_z = 0$ velocity class. When the scanned probe also reaches ω 0 ω_0, it will interrogate this same depleted class and observe the Bennett hole. The slide notes that the depth of this hole will be proportional to S 0 / 2 $S_0/2$. This is just another configuration that allows us to see a sub-Doppler feature at the line center.

Page 68:

This slide provides the visual comparison we were just discussing, showing the profoundly different results of saturating different types of lines with a standing wave. The dashed line in each plot is the unsaturated profile, and the solid blue line is the saturated profile.

Let's look at panel (a), labeled "Homogeneous Broadening." The unsaturated profile is a Lorentzian. When we apply the standing wave, the saturated profile is still a Lorentzian, but it is shorter and significantly wider due to power broadening. There is no narrow dip.

Now look at panel (b), "Inhomogeneous Broadening (Lamb Dip)." The unsaturated profile is a wide Gaussian. When this is saturated by a standing wave, the result is the same wide Gaussian but with a sharp, narrow Lamb dip carved out precisely at the center, ω 0 ω_0.

Finally, panel (c) shows the "Pump-Probe (Bennett Hole)" case. Here, a pump laser is fixed at a frequency slightly above line center, and a probe is scanned. We again see the broad Gaussian profile, but now a narrow dip—the Bennett hole—appears at the frequency of the pump laser, not necessarily at the line center.

These three plots brilliantly summarize the different spectral signatures of saturation under different conditions.

Page 69:

Up to this point, our mathematical treatment has mostly relied on rate equation models, which are a great approximation in the weak-field limit where S 0 S_0 is small. Now, we will briefly venture "Beyond the Weak-Field Approximation" to consider the case of strong saturation.

In this regime, coherence effects, which are neglected in simple rate equation models, become important. A more rigorous treatment using the density matrix formalism is required. We will not go through the derivation here, but we will present the result. The population equations yield a more complex expression for the saturated absorption coefficient.

The formula for α s (ω) $\alpha_{\mathrm{s}}(\omega)$ is given as:

α s (ω) = α 0 (ω) · γ / 2 B · 1 − 2 (ω − ω 0) A + B .

$$\alpha_{\mathrm{s}}(\omega) = \alpha_0(\omega) \cdot \frac{\gamma/2}{B} \cdot \sqrt{1 - \frac{2(\omega - \omega_0)}{A + B}}.$$

This is clearly a more complicated expression. The terms A and B, which depend on frequency and saturation, are defined on the next slide. The key point here is that for very strong fields, the shape of the dip is no longer a simple Lorentzian.

Page 70:

Here we define the terms A and B from the strong-saturation formula on the previous slide.

A is defined as the square root of, the quantity (omega minus omega naught) squared, plus (gamma over 2) squared.

A = (ω − ω 0) 2 + (γ 2) 2

$$A = \sqrt{(\omega - \omega_0)^2 + \left(\frac{\gamma}{2}\right)^2}$$

B is defined as the square root of, the quantity (omega minus omega naught) squared, plus (gamma over 2) squared times (one plus two S naught).

B = (ω − ω 0) 2 + (γ 2) 2 (1 + 2 S 0)

$$B = \sqrt{(\omega - \omega_0)^2 + \left(\frac{\gamma}{2}\right)^2 (1 + 2S_0)}$$

While the full expression is complex, it's very instructive to look at its behavior in two important limits.

First, let's consider the line-center absorption, when omega equals omega naught. In this case, the complicated formula simplifies significantly, and we find that the absorption at the dip's minimum is:

alpha sub s of omega naught equals the unsaturated value, alpha naught of omega naught, divided by the square root of (one plus two S naught).

α s (ω 0) = α 0 (ω 0) 1 + 2 S 0

$$\alpha_s(\omega_0) = \frac{\alpha_0(\omega_0)}{\sqrt{1 + 2S_0}}$$

Second, let's look at the case far off-resonance, where the absolute value of the detuning, omega minus omega naught, is much greater than the linewidth gamma. In this limit, the absorption becomes:

alpha sub s of omega equals the unsaturated value, alpha naught of omega, divided by the square root of (one plus S naught).

α s (ω) = α 0 (ω) 1 + S 0

$$\alpha_s(\omega) = \frac{\alpha_0(\omega)}{\sqrt{1 + S_0}}$$

Page 71:

Let's think about that last result. Far from the resonance, the strong-field formula for the standing wave gives us an absorption that scales as 1 1 + S 0 $\frac{1}{\sqrt{1+S_0}}$.

This should look very familiar! This is exactly the same scaling factor we derived for the saturation caused by a *single* traveling wave.

This result is perfectly consistent and makes excellent physical sense. Far off-resonance, the forward and backward waves are interacting with two completely different velocity groups that are very far apart in velocity space. The atoms resonant with the forward wave have no idea that the backward wave even exists, and vice-versa. Therefore, the total absorption is simply the sum of the absorption from two independent, single-wave saturated profiles. It's a great consistency check on our more complex formula.

Page 72:

Given that the depth and width of the Lamb dip both depend on the saturation parameter S 0 S_0, a natural question arises: what is the *optimal* value of S 0 S_0 to achieve the best-looking dip? We can define "best" as having the maximum possible contrast.

Let’s define a relative contrast function, C (S 0) $C(S_0)$, which is the absorption in the wings minus the absorption at the center, all divided by the absorption in the wings. Using our results from the strong-field analysis, this contrast is given by the expression: C of S naught equals the quantity, one over the square root of (1 plus S naught), minus one over the square root of (1 plus 2S naught), all divided by one over the square root of (1 plus

S naught). Wait, the slide has a different denominator. The denominator is 1 1 + 2 S o

. This seems to be the ratio of the depth to the central absorption. Let’s proceed with the slide’s formula.

C (S 0) = 1 1 + S 0 − 1 1 + 2 S 0 1 1 + 2 S 0

$$C(S_0) \; = \frac{\frac{1}{\sqrt{1 + S_0}} - \frac{1}{\sqrt{1 + 2\,S_0}}}{\frac{1}{\sqrt{1 + 2\,S_0}}}$$

To find the optimum value, we do what we always do in physics: we take the derivative of this function, C (S 0) $C(S_0)$, with respect to S 0 S_0, set the derivative to zero, and solve for S 0 S_0.

Performing this optimization yields the result: The optimal saturation parameter, S 0 , o p t $S_{0,\mathrm{opt}}$, is approximately equal to 1.4 1.4.

Page 73

So we've found that the optimal contrast for the Lamb dip occurs at a saturation parameter S 0 S_0 of about 1.4. What does this mean physically and practically?

At this specific value, the depth of the dip is maximized without causing excessive power broadening. It represents the sweet spot in the trade-off. If you use a lower intensity, the dip will be narrower but also much shallower and harder to see. If you use a much higher intensity, you'll broaden the dip so much that its contrast relative to the wings will actually decrease, even if the absolute absorption at the center continues to drop.

The practical implication for an experimentalist is clear: when setting up a Lamb dip experiment, you should choose your laser intensity such that the on-resonance saturation parameter, S 0 S_0, sits somewhere in the neighborhood of 1 1 to 2 2. This will give you the best combination of depth and narrowness for your signal, making it ideal for applications like frequency locking.

Page 74:

This graph visually confirms the optimization calculation we just performed. The title is "Optimal Saturation for Lamb Dip Contrast."

The plot shows the Normalized Lamb Dip Depth, which is proportional to our contrast function, on the vertical axis, versus the Saturation Parameter, S 0 S_0, on the horizontal axis. The blue curve shows how the dip depth changes as we increase the laser power. It starts at zero for S 0 = 0 $S_0 = 0$. It then rises rapidly, reaches a clear maximum, and then begins to slowly decrease for higher values of S 0 S_0.

As the annotations on the graph indicate, this maximum occurs at an optimal S 0 S_0 value of approximately 1.4. The formula shown on the graph,

Δ α ∝ 1 1 + S 0 − 1 1 + 2 S 0

$$\Delta\alpha \propto \frac{1}{\sqrt{1 + S_0}} - \frac{1}{\sqrt{1 + 2\, S_0}}$$

is proportional to the numerator of our contrast function, representing the absolute depth of the dip. This plot provides a clear, visual guide for the experimentalist: cranking up the power indefinitely is not the best strategy; there is a well-defined optimum.

Page 75:

Let's consider one more practical scenario. In a real experiment, the reflected probe wave might be much weaker than the incident pump wave due to losses at the mirror or windows. What happens if we have unequal pump and probe intensities, say the reflected intensity I 2 I_2 is much less than the incident intensity I 1 I_1?

In this case, the two counter-propagating waves are no longer symmetric. The strong incident wave acts as a powerful pump, while the weak reflected wave acts as a gentle probe. This is an intermediate case between the two-laser pump-probe setup and the ideal standing-wave Lamb dip.

The result is that the effective half-width of the observed dip changes. The new width, which we can call \(\Gamma^_\text{s}\), *is given by the* average* of the unsaturated and saturated homogeneous widths:

Γ s ∗ =γ + γ s 2

$$\Gamma_s^* = \frac{\gamma + \gamma_s}{2}$$

This makes intuitive sense. The width is an average of the hole's width (γ s γ_s) and the probe's interaction width (γ γ).

Page 76

This is a continuation of the unequal intensities case. This average width, (γ + γ s) / 2 $(\gamma + \gamma_\text{s})/2$, can be seen as the average of the width of the hole created by the strong pump and the natural width with which the weak probe interrogates it.

For the weak saturation limit, where S 0 ≪ 1 $S_0 \ll 1$, the absorption formula becomes:

α s (ω) = α 0 (ω) × dip factor

$$\alpha_\text{s}(\omega) = \alpha_0(\omega) \times \text{dip factor}$$

The dip factor is

1 − S 0 2 (a Lorentzian)

$$1 - \frac{S_0}{2}\ (a\ \text{Lorentzian})$$

The Lorentzian in this case has a width given by our new effective width, \(\Gamma^_*text{s}\). So the denominator would be \((\omega - \omega_0)^2 + (\Gamma^*_\text{s}/2)^2\).

The slide shows a slightly different form, \(\displaystyle \frac{\gamma_\text{s}^{2/4}}{(\omega - \omega_0)^2 + (\Gamma^_*text{s}/2)^2}\). These forms are related, differing by normalization constants. The key physical point is that the width of the dip is now this new \(\Gamma^*_\text{s}\).

This demonstrates a powerful experimental flexibility. By varying the relative powers of the two beams, for example by using a variable attenuator before the mirror, an experimentalist can actively tailor the shape, width, and depth of the Lamb dip for specific diagnostic purposes.

Page 77:

Alright, we have covered a great deal of ground. Let's distill it all down to the key take-aways and practical applications. This is the summary of the entire chapter.

First and foremost, the saturation of an inhomogeneous line, like a Doppler-broadened line, differs qualitatively and profoundly from the saturation of a homogeneous case. This is the single most important concept.

For a single laser beam interacting with the gas, what do we see? We see a uniform reduction of the absorption profile. There is no observable spectral hole.

To see the sub-Doppler features, we need a more sophisticated setup. We must use either a pump-probe configuration or a standing wave. These techniques are what allow us to see the Bennett holes or the Lamb dips. These are the tools that let us "dig below" the Doppler profile.

Page 78: Continuing with our key take-aways:

The width of the observed dip, be it a Bennett hole or a Lamb dip, is set not by the Doppler width, but by the power-broadened *homogeneous* width, γ s γ_s, which is equal to γ 1 + S 0 $\gamma\sqrt{1+S_0}$. This is the "sub-Doppler" nature of the technique.

The depth of the dip, for weak fields, scales roughly with the saturation parameter S 0 S_0. It grows for S 0 S_0 up to about 1, and then it saturates or even decreases in contrast for very strong fields.

These techniques are not just theoretical curiosities; they are workhorses of modern physics. They are exploited for: - Precision frequency references. The Lamb dip in molecules like iodine or acetylene, or atoms like rubidium, provides an incredibly stable and accurate frequency standard. - And, laser cooling. The velocity-selective nature of the interaction is the first step in many laser cooling schemes. By tuning a laser just below resonance, you can preferentially slow down atoms moving towards the laser, removing kinetic energy from the gas. The selection of the zero-velocity class is a key ingredient.

Page 79:

And finally, some more profound applications and a concluding thought.

Saturation spectroscopy enables fundamental tests of physics. By measuring the absolute frequencies of atomic transitions with the precision afforded by Lamb dips, we can test the predictions of quantum electrodynamics, help determine the value of fundamental constants like the fine-structure constant, and search for physics beyond the Standard Model, such as evidence for parity violation.

The final bullet point here is my concluding message to you for this chapter. Understanding every intermediate step of the physics we've discussed—from the microscopic rate equations for a single atom, to the velocity selection, to the integration over the entire ensemble, and finally to the macroscopic absorption profiles—is absolutely essential for anyone who wants to design, perform, or interpret a modern saturation spectroscopy experiment.

You now have the complete toolkit, from the ground up, to understand one of the most elegant and powerful techniques in the physicist's arsenal. Thank you.