외계행성 대기
Exoplanetary Atmospheres

Lecture 3: Flux Ratios, Phase Curves, and Radiative Transfer Fundamentals
2026 September 21 (Monday)

Based on: Sections 3.5–3.6 and Chapter 5 of Seager, *Exoplanet Atmospheres* (2010), and Chapter 2 of Heng, *Exoplanetary Atmospheres: Theoretical Concepts and Foundations* (2017)

선광일 (Kwangil Seon)
KASI / UST

Outline – Lecture 3

Why flux ratios?

An estimate of the planet-to-star flux ratio tells us whether a given planet and star can be detected with a given telescope, at which wavelengths, and with what precision. Last lecture gave us the ingredients: surface fluxes, the flux at Earth, temperatures, and albedos. We now assemble them.

- Assume star and planet radiate as blackbodies (a first estimate; brightness temperatures refine it).
- A planet has **two spectral peaks**: a short-wavelength peak of *scattered starlight*, at the same wavelength as the star, and a long-wavelength peak of *thermal emission* at the planet's own temperature.
- During transit a third quantity appears: the *transmission* flux ratio, the fraction of starlight blocked by the atmospheric annulus.

The planet flux at Earth is $F_{\oplus,p} = F_{S,p}(R_p/D_{\oplus})^2$ and the stellar flux at Earth is $F_{\oplus,*} = F_{S,*}(R_*/D_{\oplus})^2$, so the distance cancels:

$$\boxed{\frac{F_{\oplus,p}}{F_{\oplus,*}} = \frac{F_{S,p}\,R_p^2}{F_{S,*}\,R_*^2}}$$

valid for total fluxes and for frequency-dependent fluxes alike.

Spectra of the Solar System seen from 33 light years

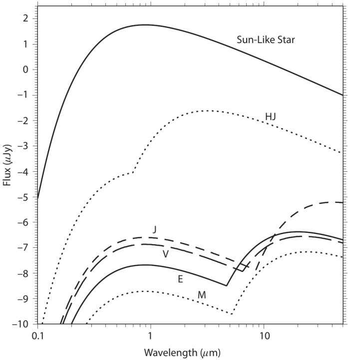

Figure 1: Approximate blackbody spectra (in 10^{-26} W m^{-2} Hz^{-1}) of the Sun (5750 K), Jupiter, Venus, Earth, Mars, and a representative hot Jupiter, seen from 33 light years (10 pc). Reproduced from Seager (2010), Fig. 3.8.

- Each planet shows the two peaks: reflected light near 0.5 μm (shape of the stellar spectrum, scaled by $A_g R_p^2/a^2$) and thermal emission peaking at $\lambda_{\max} \simeq 2898\,\mu\mathrm{m\,K}/T$: $\sim 10\,\mu$m for Earth and Venus, $\sim 25\,\mu$m for Jupiter, $\sim 2\,\mu$m for the hot Jupiter.
- The planet-star contrast is large at *all* wavelengths, but the thermal peak of a hot Jupiter reaches within 10^{-3} of the star, the Earth's thermal peak within 10^{-7}, while in the visible the contrast is 10^{-9}–10^{-10}.
- Venus is brighter than Earth in reflected light (high albedo, closer to the Sun) but similar in the thermal infrared: its cloud tops radiate at $T_{\mathrm{eff}} \simeq 230$ K.
- This high contrast is *the* primary challenge for detecting radiation from exoplanets.

Thermal emission flux ratio

On the illuminated side, equate the planet's thermal emission with the absorbed stellar radiation — the energy balance of last lecture with the redistribution factor f, $4\pi R_p^2 F_{\mathrm{S},p}/f = (1-A_\mathrm{B})F_{\mathrm{S},*}(R_*/a)^2\pi R_p^2$ — to get the *total* flux ratio

$$\boxed{\frac{F_{\oplus,p}}{F_{\oplus,*}} = \frac{F_{\mathrm{S},p}R_p^2}{F_{\mathrm{S},*}R_*^2} = \left(\frac{R_p}{a}\right)^2 \frac{f}{4}(1-A_\mathrm{B})}$$

- Here f is the correction factor to 4π that describes the unknown redistribution of the absorbed stellar radiation, and the fluxes are total fluxes, integrated over all frequencies.
- The ratio depends on R_p/a, not on R_p/R_*: the planet reradiates the fraction $\pi R_p^2/4\pi a^2$ of the stellar luminosity it intercepts, times $(1-A_\mathrm{B})$, times the redistribution factor.
- For a hot Jupiter at $a = 0.05$ AU with $A_\mathrm{B} = 0.1$ and $f = 1$: $(R_\mathrm{J}/a)^2 \times 0.9/4 \simeq 2\times10^{-5}$ bolometrically. Yet the observed ratio in an infrared band is $\sim 10^{-3}$, because the planet's emission is concentrated at long wavelengths where the star is faint — which is why we need a frequency-dependent estimate.

Thermal emission flux ratio at a given frequency

A more useful planet-star flux ratio estimate is **frequency dependent**, because telescope observations are limited to a fixed frequency or wavelength range. Approximate the star and planet surface fluxes by blackbody fluxes $\pi B(T,\nu)$, with T the effective temperature of the star and the equilibrium temperature of the planet:

$$\boxed{\frac{F_{\oplus,p}(\nu)}{F_{\oplus,*}(\nu)} = \frac{F_{\mathrm{S},p}(\nu)R_p^2}{F_{\mathrm{S},*}(\nu)R_*^2} = \frac{\mathrm{e}^{h\nu/k_\mathrm{B}T_{\mathrm{eff},*}}-1}{\mathrm{e}^{h\nu/k_\mathrm{B}T_{\mathrm{eq},p}}-1}\,\frac{R_p^2}{R_*^2}}$$

- The ν^3 prefactors of the two Planck functions cancel; only the exponential factors and the area ratio remain.
- In the Wien tail of the planet ($h\nu \gg k_\mathrm{B}T_\mathrm{eq}$) the ratio is $\simeq \mathrm{e}^{-h\nu(1/k_\mathrm{B}T_\mathrm{eq}-1/k_\mathrm{B}T_{\mathrm{eff},*})}(R_p/R_*)^2$, exponentially small; it rises steeply toward the infrared and saturates in the Rayleigh–Jeans regime (next slide).
- Note that it is now R_p/R_* that matters, not R_p/a: at a fixed planet temperature, a smaller star gives a larger contrast. This is the same “M-dwarf opportunity” as for transits.
- Example: HD 189733b at 8 μm with $R_p/R_* = 0.155$, $T_\mathrm{eq} = 1200\,\mathrm{K}$, $T_{\mathrm{eff},*} = 5000\,\mathrm{K}$ gives $(\mathrm{e}^{3.6}-1)/(\mathrm{e}^{1.5}-1)\times 0.024 = 0.0034$, close to the observed eclipse depth of 0.34% (Lecture 1).

The Rayleigh–Jeans limit

At mid-infrared wavelengths a further simplification applies: for reasonable planet and star temperatures both blackbody fluxes are in the **Rayleigh–Jeans** region, $h\nu/k_\mathrm{B}T \ll 1$. With $\mathrm{e}^x - 1 \simeq x$ the Planck function becomes

$$\pi B(T,\nu) \simeq \pi \frac{2\nu^2}{c^2} k_\mathrm{B} T,$$

linear in T, and the flux ratio becomes

$$\boxed{\frac{F_{\oplus,p}(\nu)}{F_{\oplus,*}(\nu)} = \frac{F_{\mathrm{S},p}(\nu)R_p^2}{F_{\mathrm{S},*}(\nu)R_*^2} \simeq \frac{T_{\mathrm{eq},p}}{T_{\mathrm{eff},*}}\left(\frac{R_p}{R_*}\right)^2}$$

in the frequency range where the approximation is valid.

- This is the *maximum* contrast the thermal emission can reach: the temperature ratio times the area ratio.
- **Hot Jupiter:** $(1500/5800) \times (0.1)^2 \simeq 2.6 \times 10^{-3}$, approached already at $\lambda \gtrsim 10\,\mu\mathrm{m}$ (for $T_p = 1500\,\mathrm{K}$, $h\nu/k_\mathrm{B}T = 1$ at $\lambda = 9.6\,\mu\mathrm{m}$).
- **Earth:** $(255/5800) \times (0.0092)^2 \simeq 3.7 \times 10^{-6}$, reached only at $\lambda \gtrsim 100\,\mu\mathrm{m}$; at $10\,\mu\mathrm{m}$ the actual ratio is $\sim 10^{-7}$ because Earth is still on the Wien side.
- The Rayleigh–Jeans form makes explicit why the infrared is the place to detect thermal emission: the contrast saturates at T_p/T_* instead of decaying exponentially.

Using brightness temperatures instead of blackbodies

Real planets and stars deviate from blackbodies, sometimes strongly (the M-dwarf spectrum of last lecture, molecular bands in the planet). The honest version of the Rayleigh–Jeans flux ratio uses the **brightness temperatures** at each frequency:

$$\boxed{\frac{F_{\oplus,p}(\nu)}{F_{\oplus,*}(\nu)} = \frac{F_{S,p}(\nu)R_p^2}{F_{S,*}(\nu)R_*^2} = \frac{T_{b,p}(\nu)}{T_{b,*}(\nu)}\left(\frac{R_p}{R_*}\right)^2} \qquad \text{(Rayleigh–Jeans regime)},$$

and, outside the Rayleigh–Jeans regime, the ratio of Planck functions evaluated at $T_{b,p}(\nu)$ and $T_{b,*}(\nu)$.

- This is exactly how a measured **secondary-eclipse depth** is turned into a dayside brightness temperature: with R_p/R_* from the transit and a stellar model for $T_{b,*}(\nu)$, invert the Planck function for $T_{b,p}(\nu)$.
- The stellar brightness temperature matters: at 3.6–8 μm a K or M dwarf has $T_{b,*}$ well below its $T_{\rm eff}$ because of molecular absorption, which raises the apparent planet-to-star contrast.
- $T_{b,p}(\nu)$ across several bands maps out the emission spectrum: bands with lower T_b are absorption features formed high in a (normally stratified) atmosphere; bands with higher T_b indicate a temperature inversion.

Scattered-light flux ratio at full phase

Scattered starlight is typically at visible wavelengths (extending to longer wavelengths for colder planets). At $\alpha = 0$, start again from the thermal flux ratio and replace the absorbed-and-reradiated factor by a scattered one:

$$\frac{f}{4}(1 - A_{\rm B}) \quad \longrightarrow \quad \frac{g}{4}A_{\rm S}(\nu),$$

where $A_{\rm S}(\nu)$ is the spherical albedo (we want a wavelength-dependent ratio; monochromatic scattering assumed) and g, like f, is a correction to 4π: the solid angle into which the incident radiation is scattered compared to 4π. Scattered light cannot be advected around the planet; it leaves the illuminated hemisphere, with a directional pattern. Thus

$$\frac{F_{\oplus,p}(\nu)}{F_{\oplus,*}(\nu)} = \frac{R_p^2}{a^2}\frac{g}{4}A_{\rm S}(\nu).$$

The critical step is g: the ratio of energy emitted at zero phase angle to that emitted at all phase angles. With $E_{\rm scat}$ from the spherical-albedo derivation and Ψ_α the fraction scattered toward the observer,

$$\frac{g}{4\pi} = \frac{E_{{\rm scat},\alpha=0}}{E_{\rm scat}} = \frac{\Psi_{\alpha=0}(\nu)}{2\pi\int_0^\pi \Psi_\alpha(\nu)\sin\alpha\,{\rm d}\alpha} = \frac{1}{\pi q(\nu)} = \frac{A_g(\nu)}{\pi A_{\rm S}(\nu)} \quad \Longrightarrow \quad \boxed{g = \frac{4A_g(\nu)}{A_{\rm S}(\nu)}}$$

using $\Psi_{\alpha=0} = A_g/\pi$, the phase function Φ_α, and the phase integral $q = A_{\rm S}/A_g$ (last lecture).

Scattered-light flux ratio: the result and numbers

Substituting g, the spherical albedo cancels and the full-phase visible flux ratio is

$$\boxed{\frac{F_{\oplus,p}(\nu)}{F_{\oplus,*}(\nu)} = A_g(\nu)\,\frac{R_p^2}{a^2}}$$

- A common misunderstanding: the visible flux ratio at full phase involves the *geometric* albedo, which describes the light scattered toward the observer at $\alpha = 0$, not the spherical or Bond albedo, which describe the energy scattered in all directions. That is precisely what g converts.
- The ratio falls as a^{-2} and does not involve R_* at all (for a fixed stellar flux at the planet, a bigger star only means a bigger a for the same temperature).
- **Jupiter:** $A_g(500\,\mathrm{nm}) = 0.5$, $a = 5.2$ AU, $R_p = R_\mathrm{J}$ gives $F_{\oplus,p}/F_{\oplus,*} = 4.4 \times 10^{-9}$.
- **Scaled to 0.05 AU:** 10^4 times higher, 4.4×10^{-4}. The actual hot Jupiters are darker: Jupiter is bright because of its icy NH_3 and H_2O clouds, whereas hot Jupiters lack bright clouds and have strong Na/K absorption, so $A_g \lesssim 0.1$ and the ratio is $\lesssim 10^{-4}$ — consistent with the Kepler and MOST upper limits and detections.
- **Earth:** $A_g(500\text{–}800\,\mathrm{nm}) = 0.2$, $a = 1$ AU: 2.1×10^{-10} at $\alpha = 0$ (the number quoted in Lecture 1).

Transmission flux ratio

During transit, starlight passes through the planet atmosphere and picks up its spectral features: the **transmission spectrum**. What is measured is the total (star + planet) flux in transit divided by the out-of-transit flux. Let R_p be the radius where the planet is opaque at all wavelengths of interest and A_H the additional radial extent of the atmosphere. Splitting the stellar disk into the unobscured part, the opaque planet ($I_p = 0$), and the atmospheric annulus:

$$F_{\mathrm{in\,trans},\oplus}(\nu) = \int_0^{2\pi} \int_{(R_p+A_H)/D_\oplus}^{R_*/D_\oplus} I_* \cos\omega \sin\omega \,\mathrm{d}\omega\,\mathrm{d}\phi + \int_0^{2\pi} \int_{R_p/D_\oplus}^{(R_p+A_H)/D_\oplus} I_{\mathrm{atm}} \cos\omega \sin\omega \,\mathrm{d}\omega\,\mathrm{d}\phi .$$

Dividing by the out-of-transit stellar flux $F_{*,\oplus} = F_{\mathrm{S},*}(R_*/D_\oplus)^2$ and assuming $A_H \ll R_p$,

$$\boxed{\frac{F_{\mathrm{in\,trans},\oplus}(\nu)}{F_{*,\oplus}(\nu)} = 1 - \left(\frac{R_p}{R_*}\right)^2 - \frac{2R_pA_H}{R_*^2}\left[1 - \frac{F_{\mathrm{S},p,\mathrm{trans}}(\nu)}{F_{\mathrm{S},*,\mathrm{trans}}(\nu)}\right]}$$

- First term: the star; second: the opaque disk (the transit depth); third: the annulus $2\pi R_pA_H$ over the stellar disk, times the fraction of starlight the annulus *blocks*. $F_{\mathrm{S},p,\mathrm{trans}}/F_{\mathrm{S},*,\mathrm{trans}}$ is the transmitted fraction, computed in Chapter 6 from the slant optical depth.

How large is the transmission signal?

The maximum signal occurs where the annulus is completely opaque ($F_{S,p,\mathrm{trans}} = 0$):

$$\boxed{\frac{F_{p,\oplus,\mathrm{trans}}(\nu)}{F_{*,\oplus}(\nu)} = \frac{2R_pA_H}{R_*^2}\left[1 - \frac{F_{S,p,\mathrm{trans}}}{F_{S,*,\mathrm{trans}}}\right] \approx \frac{2R_pA_H}{R_*^2}}$$

- The choice of A_H does not matter as long as it is large enough: above some height there is no atmosphere left to attenuate, the transmitted fraction is 1, and the bracket vanishes. What matters is the height at which the slant path becomes opaque, a few scale heights above the opaque radius.
- **Hot Jupiter** around a Sun-sized star: an H/He atmosphere with $T_{\mathrm{eq}} \simeq 1000\,\mathrm{K}$ has a scale height $H \simeq 500\,\mathrm{km}$ (Chapter 9). Taking the atmosphere to extend $A_H = 5H$: $2R_pA_H/R_*^2 = 2 \times (7.1 \times 10^4) \times (2.5 \times 10^3)/(7 \times 10^5)^2 \simeq 7 \times 10^{-4}$, i.e. of order 10^{-3} (Seager quotes 1.5×10^{-3}).
- One part per thousand may seem very small. In fact it is a *large* number in the field of exoplanet atmosphere detection: compare with the $\lesssim 10^{-4}$ reflected-light and $\sim 10^{-3}$ thermal ratios above, and note that it requires no separation of the planet's light from the star's, only precise photometry in and out of transit.

Transmission signals of cooler planets, and the three channels

- **Jupiter analog** (a Jupiter-sized planet in a Jupiter-like orbit, $T_{\rm eq} \simeq 120\,\mathrm{K}$, $H = 24\,\mathrm{km}$): the estimated transmission strength is only ~ 4–7×10^{-5}. Cold planets have small scale heights and thus weak transmission spectra, even though their transit depth is the same.
- **Earth twin** ($H = 8\,\mathrm{km}$, N_2 atmosphere): $\sim 2 \times 10^{-7}$ (Lecture 1). This is why terrestrial-planet transmission spectroscopy targets M dwarfs, where R_*^2 is 25–100 times smaller.
- The transmission technique has been very successful for hot Jupiters around bright stars: Na, K, H_2O, CH_4, CO, CO_2, He, and escaping H have all been identified this way, and JWST now routinely delivers transmission spectra of sub-Neptunes.

The three observational channels and what controls each:

channel	flux ratio	key planet property	geometric factor
thermal emission	$(T_{\mathrm{b},p}/T_{\mathrm{b},*})(R_p/R_*)^2$	temperature $T_{\mathrm{b}}(\nu)$	$(R_p/R_*)^2$
reflected light	$A_g R_p^2/a^2$	geometric albedo $A_g(\nu)$	$(R_p/a)^2$
transmission	$2R_pA_H/R_*^2$	scale height H	R_p/R_*^2

Each channel favors different planets: emission and transmission favor hot, large planets around small stars; reflection favors close-in planets but is limited by their low albedos.

Visible-wavelength phase curves

An exoplanet goes through illumination phases like the Moon. We want the flux ratio not only at $\alpha = 0$ but as a function of α. Recall the normalized phase function

$$\Phi_\alpha(\nu) = \frac{\Psi_\alpha(\nu)}{\Psi_{\alpha=0}(\nu)}, \qquad \Psi_\alpha(\nu) = \int_{\alpha-\pi/2}^{\pi/2} \int_{-\pi/2}^{\pi/2} I_{\text{inc}}(\theta', \phi', \nu)\, p \cos(\alpha - \phi) \cos\phi \cos^3\theta \, \mathrm{d}\theta \, \mathrm{d}\phi .$$

The Lambert sphere is the one case with an analytic phase curve. With $p = 1$ and uniform incident intensity, the θ integral gives $4/3$ and the ϕ integral gives $\int_{\alpha-\pi/2}^{\pi/2} \cos(\alpha - \phi) \cos\phi \, \mathrm{d}\phi = \frac{1}{2}[\sin\alpha + (\pi - \alpha)\cos\alpha]$, so after normalizing to $\alpha = 0$ (where the integral is $\pi/2$),

$$\boxed{\Phi_\alpha = \frac{1}{\pi}\left[\sin\alpha + (\pi - \alpha)\cos\alpha\right]}$$

frequency independent, with $\Phi_0 = 1$, $\Phi_{\pi/2} = 1/\pi$, $\Phi_\pi = 0$.

- At quadrature ($\alpha = 90^\circ$) a Lambert sphere is $1/\pi \simeq 0.32$ as bright as at full phase, *less than half* (exercise 3.7): the visible illuminated half-disk is lit obliquely near the terminator and its brightest region (the substellar point) sits at the limb, seen at grazing angle.
- The phase integral follows: $q = 2\int_0^\pi \Phi_\alpha \sin\alpha \, \mathrm{d}\alpha = 3/2$.

Phase curves of Solar System bodies

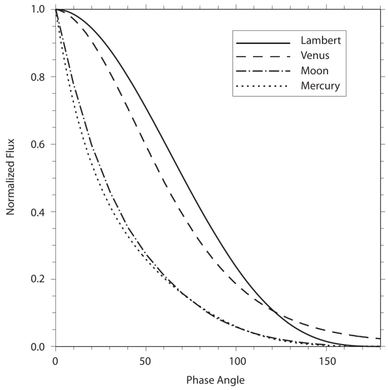

Figure 2: Visible-wavelength phase curves of Venus, Mercury, and the Moon, with the Lambert sphere (rightmost solid line). From formulae in Cox (2000); reproduced from Seager (2010), Fig. 3.9.

- **Venus**, with its thick cloud deck, follows the Lambert curve closely out to $\alpha \sim 100^\circ$ and then falls faster (forward scattering by cloud droplets brightens it again at very large α, not shown).
- The **Moon** and **Mercury** — airless, rough, regolith-covered — fall off much faster: their surfaces backscatter (the *opposition surge* near $\alpha = 0$) and shadowing in the rough surface darkens the oblique views. At quadrature the Moon is only ~ 0.1 of its full-phase brightness, not 0.32.
- Lesson: the phase curve shape encodes the *scattering physics* (clouds versus bare surface, particle size, surface roughness). For exoplanets it is a second observable beyond A_g.

Flux ratio at all phases and an Earth example

The normalized phase curve combined with the full-phase flux ratio gives the planet-star flux ratio at any phase:

$$\boxed{\frac{F_{p,\oplus,\alpha}(\nu)}{F_{*,\oplus}(\nu)} = A_g(\nu)\,\frac{R_p^2}{a^2}\,\Phi_\alpha(\nu)}$$

Earth orbiting the Sun as an exoplanet.

- At $\alpha = 0$ with $A_g(500\text{–}800\,\text{nm}) = 0.2$, $a = 1$ AU, $R_p = R_\oplus$: 2.1×10^{-10}. Earth is almost ten billion times fainter than the Sun — and $\alpha = 0$ is when the planet is directly behind the star.
- At **maximum elongation** ($\alpha = 90^\circ$), the natural configuration for direct imaging, approximating Earth as a Lambert sphere: $\Phi_{90^\circ} = 1/\pi$ and the ratio is 1.2×10^{-10}, a factor π lower. (Earth's real phase curve is not Lambertian: ocean glint and forward-scattering clouds modify it.)
- For an edge-on orbit the full range $0 < \alpha < 180^\circ$ is swept once per orbit; the amplitude of the reflected-light phase curve is $A_g R_p^2/a^2$ and its *shape* constrains Φ_α, hence q and A_S. For inclination i, the accessible range is $90^\circ - i \le \alpha \le 90^\circ + i$ (exercise 3.8).

Infrared phase curves

Phase variation also occurs at *infrared* wavelengths, from thermal emission, and not just at visible wavelengths from scattered starlight. Thermal phase variation arises when the emission is dominated by reemitted absorbed starlight rather than by internal energy.

- The prime examples are the **tidally locked** (synchronously rotating) hot Jupiters, hot Neptunes, and hot Earths with $a < 0.05$ AU, strongly heated by their stars. They present the same face to the star at all times, just as the Moon does to Earth.
- The permanent day and night sides set up a temperature gradient that atmospheric circulation may or may not even out. How the absorbed stellar energy is circulated is a complex topic (Chapter 10).

Simplest case: a hot tidally locked planet where the incident stellar energy is absorbed in one layer, instantaneously heats that surface element, and is reemitted, with no advection. Thermal emission is isotropic, so each element reradiates like a Lambert scatterer and the problem is solved in analogy with the Lambert sphere: the normalized thermal phase curve is again

$$\Phi_\alpha = \frac{1}{\pi}\left[\sin\alpha + (\pi-\alpha)\cos\alpha\right],$$

peaking at secondary eclipse and vanishing at transit (no nightside emission at all), with the dayside temperature corresponding to $f' = 2/3$.

Infrared phase curves: the general case

Beyond the simplest case, if the thermal flux distribution over the planet's surface is available from a model atmosphere or a circulation model, the thermal phase curve is still computed from the planet-star flux ratio, with the phase-dependent planetary flux

$$F_{\alpha,\mathrm{thermal}}(\nu) = \int_{-\pi/2}^{\pi/2}\int_{-\pi/2}^{\pi/2} I_{\alpha,\mathrm{thermal}}(\theta,\phi,\nu)\cos\phi\cos^2\theta\,\mathrm{d}\theta\,\mathrm{d}\phi,$$

integrated over the hemisphere facing the observer at phase α (the subscript "thermal" denotes the thermal-infrared component of the intensity), and the flux ratio follows with R_p^2/R_*^2 and the stellar flux at that frequency.

- The shape of $I_{\alpha,\mathrm{thermal}}(\theta,\phi)$ encodes the longitudinal temperature map: a hot spot shifted east of the substellar point produces a phase-curve peak *before* secondary eclipse (Lecture 1).
- The nightside flux sets the minimum of the phase curve near transit; the day–night amplitude measures how efficiently heat is redistributed, i.e. the effective f'.
- Real hot Jupiters (HD 189733b, HD 209458b) show amplitudes reduced relative to the no-redistribution Lambert curve and offset peaks: heat is carried to the nightside by winds. Ultra-hot Jupiters and planets on eccentric orbits (HD 80606b) show the full range from strong to weak redistribution.

Heng's approach: start with the optical depth

"It is possible to sit through an entire radiative transfer class and not develop an intuition for the subject at all" (Heng). Rather than derive the transfer equation first, we introduce the one quantity needed to understand it: the **optical depth** τ.

- It quantifies whether a medium is transparent or opaque, and therefore which approximation describes the passage of radiation.
- It explains why the size of a star or a planet varies with wavelength, identifies the surface from which photons reach our telescopes, and defines the transit radius.

Clouds in the sky: some large clouds are transparent, some small clouds are opaque. Physical size alone is a poor indicator; two more ingredients matter: how tightly the material is packed (the number density n) and how absorbent each particle is (the cross section σ):

$$\boxed{\tau \equiv \int n\,\sigma\,\mathrm{d}x}$$

a macroscopic quantity times a microscopic one, integrated along the path. A photon travels, on average, until it has traversed $\boxed{\tau \sim 1}$: τ is roughly the number of interactions the photon suffers. Transparent clouds have $\tau \ll 1$ (no interaction on average); opaque clouds $\tau \gg 1$ (many absorptions and reemissions). Walking through fog or snow, visibility extends to $\tau \sim 1$.

What $\tau \sim 1$ explains

- **Why the Sun is yellow, not an X-ray source.** The core is at 10^6–10^7 K (X-rays by Wien's law), but it sits at $\tau \gg 1$. What we see is the *photosphere*, the surface where $\tau \sim 1$ as measured from outside: the surface of last absorption or scattering before the photons stream to us.
- **Why the Sun has different sizes.** Its "radius" is the $\tau \sim 1$ surface, and since τ is wavelength dependent, the Sun observed in X-rays, UV, visible, or radio has different radii. (Styrofoam is opaque in the visible but transparent to microwaves.)
- **How deep starlight penetrates a hot Jupiter.** To the $\tau \sim 1$ surface at that wavelength: the concept applies symmetrically to emission from an object and absorption by it.
- **The transit radius.** During transit the measured radius picks out the chord across the limb where the *chord* optical depth is ~ 1 (next slide).
- **Which regime.** For $\tau \ll 1$ radiation is freely streaming and propagates at c; for $\tau \gg 1$ it is absorbed and reemitted many times over a short distance and its passage resembles *diffusion*. The two-stream approximation (Heng Chapter 3) is designed for $\tau \lesssim 1$; the diffusion approximation for $\tau \gg 1$.

Photospheric radius versus transit radius

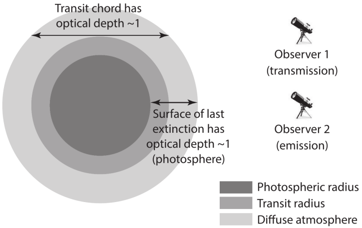

Figure 3: Observer 1 measures the transit radius (transit chord with $\tau \sim 1$); observer 2 sees photons from the photosphere, the surface of last extinction with $\tau \sim 1$ toward the observer. The transit radius exceeds the photospheric radius, exaggerated here. Reproduced from Heng (2017), Fig. 2.1.

Two observers, two radii, both defined by $\tau \sim 1$:

- **Emission (observer 2):** the photosphere is where the *vertical* optical depth from the top is ~ 1.
- **Transmission (observer 1):** the transit radius is where the *chord* (slant) optical depth is ~ 1.

For an isothermal hydrostatic atmosphere with scale height H (Heng, problem 2.8.6), along a chord at impact height $z = 0$ with $x^2 \simeq 2Rz$,

$$\tau_{\rm chord} = n_{\rm ref}\,\sigma\,\sqrt{2\pi HR}, \qquad \tau_{\rm vert} = n_{\rm ref}\,\sigma\,H,$$

so $\tau_{\rm chord}/\tau_{\rm vert} = \sqrt{2\pi R/H}$: for a hot Jupiter ($R/H \sim 150$) about 30. The chord reaches $\tau \sim 1$ at lower density, i.e. *higher* altitude, by $\Delta z = H \ln\sqrt{2\pi R/H} \simeq 3.4H$. Hence the transit radius is larger than the photospheric radius, and transmission spectra probe lower pressures than emission spectra.

The random walk of a photon (Heng, problem 2.8.2)

Photons created at the center of the Sun must travel $R_\odot \simeq 7\times 10^5$ km to escape, moving a mean free path $l_{\rm mfp}$ between interactions.

- **Transparent Sun:** escape time $R_\odot/c \simeq 2.3$ s.
- **Mildly opaque** ($1 \lesssim \tau \lesssim 10$): the number of interactions is $N \sim R_\odot/l_{\rm mfp} \sim \tau$.
- **Very opaque:** the photon random-walks. Its mean displacement is zero, but the mean *square* displacement after N steps is $Nl_{\rm mfp}^2$. Setting $\sqrt{N}l_{\rm mfp} = R_\odot$ gives $N \sim (R_\odot/l_{\rm mfp})^2 = \tau^2$.
- With $n \sim 10^{24}\,{\rm cm}^{-3}$ and Thomson scattering $\sigma \sim 10^{-24}\,{\rm cm}^2$: $l_{\rm mfp} \sim 1$ cm, $\tau \sim 7\times 10^{10}$, $N \sim 5\times 10^{21}$, and the escape time $Nl_{\rm mfp}/c \sim 10^{11}$ s $\sim$ a few thousand years. (The real answer, 10^4–10^5 yr, differs because the density is far from uniform and absorption plus reemission is not a pure random walk — the key weakness of the estimate.)

The same reasoning tells us that in an optically thick atmosphere radiative energy transport is *diffusive*, with an effective diffusion coefficient $\sim c\,l_{\rm mfp}/3$, which is the basis of the Rosseland mean opacity below.

Three ways to write the optical depth

There are three measures of how absorbent a medium is, hence three equivalent definitions of τ. *Extinction* means absorption plus scattering.

quantity	symbol	units (cgs)	optical depth
cross section (per particle)	σ	cm^2	$\tau = \int n\sigma\,\mathrm{d}x$
opacity (cross section per unit mass)	κ	$\mathrm{cm}^2\,\mathrm{g}^{-1}$	$\tau = \int \kappa\,\mathrm{d}\tilde{m} = \int \kappa\rho\,\mathrm{d}x$
extinction coefficient (per length)	α_e ($\equiv \kappa$ in Seager)	cm^{-1}	$\tau = \int \alpha_e\,\mathrm{d}x$

- $\tilde{m}$ is the **column mass** (mass per unit area). In an atmosphere the pressure is the weight of the overlying column, $P = \tilde{m}g$ (Newton's second law per unit area), so $\tau = \int \kappa\,\mathrm{d}P/g$: for constant κ, $\tau = \kappa P/g$, which is why pressure is the natural vertical coordinate of atmospheric radiative transfer.
- The cross section may be visualized as the target area presented by an atom, molecule, or particle.
- The reciprocal of the extinction coefficient is the **mean free path**, $l_{\mathrm{mfp}} = 1/\alpha_e = 1/n\sigma$.
- **Terminology warning.** Some references call σ, κ, and α_e collectively the "extinction coefficient" although their units differ. Seager uses κ for the coefficient in m^{-1} and κ_m for the mass-independent opacity in $\mathrm{m}^2\,\mathrm{kg}^{-1}$; Heng uses α_e and κ respectively. Always check the units.

Opacity: the most important atmospheric property

The most important physical quantity in the atmosphere that affects the transfer of radiation is the **opacity**: how opaque a substance is, how hard it is for radiation to pass through it.

- Opacity depends on the number density of particles and on their absorbing, scattering, and emitting properties, which in turn depend on temperature, pressure, and frequency.
- To describe radiative transfer we must therefore define coefficients for the absorption and emission of radiation. The word “opacity” refers to the absorption and scattering (i.e. extinction) coefficients of the atmospheric particles.
- For now the underlying microscopic processes of photon absorption and emission are hidden in these macroscopic coefficients; Chapter 8 (and Heng Chapter 5) supply the microscopic details: line lists, line profiles, Rayleigh scattering, condensates.
- Many different symbols are used for the opacity coefficients in other textbooks and in the literature (see the table of three definitions above); keep track of the units.

The **monochromatic extinction coefficient** $\kappa(\mathbf{x},\nu)$ [m^{-1}] is defined by the energy removed from a beam of radiation in a volume $\mathrm{d}V$ and solid angle $\mathrm{d}\Omega$, per unit time and frequency:

$$\boxed{\mathrm{d}E = \kappa(\mathbf{x},\nu)\, I(\mathbf{x},\hat{\mathbf{n}},\nu)\, \mathrm{d}V\, \mathrm{d}\Omega\, \mathrm{d}\nu\, \mathrm{d}t}$$

Absorption and scattering

κ includes all processes that remove energy from the beam:

$$\boxed{\kappa(\mathbf{x},\nu) = \alpha(\mathbf{x},\nu) + \sigma_s(\mathbf{x},\nu)}$$

- **True absorption** α: the photon is destroyed. For example, a photon is absorbed by a molecule, and the excitation energy is then converted to kinetic energy of the gas by a subsequent collision. The photon's energy has joined the thermal pool.
- **Scattering** σ_s: the photon is removed from the beam by a change of direction (and possibly a change of energy). The subscript s distinguishes the scattering coefficient σ_s [m^{-1}] from the cross section σ [m^2]. Examples: Rayleigh scattering by molecules, Mie scattering by cloud particles.
- Extinction is *isotropic* in a static medium: the absorbing or scattering particle does not care from which direction the photon comes, so κ has no $\hat{\mathbf{n}}$ dependence. (In a moving medium the Doppler shift introduces an angle dependence, which we ignore for atmospheres.)
- The distinction between α and σ_s is physical, not bookkeeping: absorbed energy heats the gas and is reemitted thermally at the local temperature; scattered energy keeps its frequency and carries no information about the local temperature.

The emission coefficient: thermal emission and scattering

The **monochromatic emission coefficient** $\varepsilon(\mathbf{x}, \hat{\mathbf{n}}, \nu)$ [J m^{-3} sr^{-1} s^{-1} Hz^{-1}] describes the energy emitted into a volume dV within a solid angle dΩ per unit time and frequency,

$$\mathrm{d}E = \varepsilon(\mathbf{x}, \hat{\mathbf{n}}, \nu)\,\mathrm{d}V\,\mathrm{d}\Omega\,\mathrm{d}t\,\mathrm{d}\nu .$$

Planetary-atmosphere emission includes both thermal emission and scattering, because both processes add to the beam.

- **Thermal emission**, by the Kirchhoff–Planck relation (*Kirchhoff's law*, valid in local thermodynamic equilibrium, Section 5.4):

  $$\boxed{\varepsilon_{\mathrm{therm}}(\mathbf{x}, \nu) = \alpha(\mathbf{x}, \nu)\,B(\mathbf{x}, \nu)}$$

  where B is the Planck function at the local temperature. Thermal emission is isotropic in a static medium: atoms and molecules have no preferred direction of radiation.
- **Scattering** of radiation arriving from all directions $\hat{\mathbf{n}}'$ into the beam direction $\hat{\mathbf{n}}$:

  $$\boxed{\varepsilon_{\mathrm{scat}}(\mathbf{x}, \hat{\mathbf{n}}, \nu) = \sigma_s(\mathbf{x}, \nu)\,\frac{1}{4\pi}\oint_{\Omega'} P(\Theta)\,I(\mathbf{x}, \hat{\mathbf{n}}', \nu)\,\mathrm{d}\Omega'}$$

  with the scattering angle $\cos\Theta = \hat{\mathbf{n}} \cdot \hat{\mathbf{n}}'$; primes denote the direction of incidence. We consider only *coherent* scattering, in which the photon retains its frequency.

The single-scattering phase function

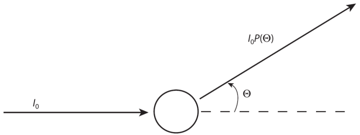

Figure 4: The scattering angle Θ used in $P(\Theta)$: incident intensity I_0, scattered intensity $I_0P(\Theta)$. Seager (2010), Fig. 5.1.

Scattering can be anisotropic, so we define the dimensionless **single-scattering phase function** $P(\Theta)$: the redirection of the incident intensity into the outgoing intensity, i.e. the 3D directional scattering probability. It is normalized so that

$$\boxed{\frac{1}{4\pi}\oint_{\Omega} P(\Theta)\,\mathrm{d}\Omega = 1}$$

- If the scattering is isotropic, $P(\Theta) = 1$ and the scattering emissivity reduces to $\varepsilon_{\text{scat}}(\mathbf{x},\nu) = \sigma_s(\mathbf{x},\nu)\,J(\mathbf{x},\nu)$, with J the mean intensity of Chapter 2: an isotropic scatterer simply reemits the mean intensity it sees.
- Θ is the angle between the incident direction $\hat{\mathbf{n}}'$ and the scattered direction $\hat{\mathbf{n}}$, $\cos\Theta = \hat{\mathbf{n}}\cdot\hat{\mathbf{n}}'$; for a spherically symmetric particle P depends only on Θ, not on the absolute orientation.

Phase functions in practice

- **Rayleigh scattering** by molecules has $P(\Theta) = \frac{3}{4}(1+\cos^2\Theta)$: symmetric between forward and backward, mildly anisotropic (Chapter 8).
- **Cloud and haze particles** comparable to or larger than the wavelength have strongly forward-peaked P (Mie theory, Chapter 8), often approximated by the Henyey–Greenstein function with asymmetry parameter $g_0 = \langle\cos\Theta\rangle$; the two-stream treatment of Heng Chapter 3 uses only g_0.
- The origin of the phase function is microscopic: the angular pattern of dipole radiation for Rayleigh scattering, and the interference of waves scattered from different parts of a large particle for Mie scattering.
- P must not be confused with the *illumination* phase function Φ_α used for phase curves (Section 3.6): both are called "phase function" for historical reasons, but P describes a single scattering event and Φ_α the brightness of the whole planet versus orbital phase. The relation between the two, through multiple scattering in the atmosphere, is what a reflected-light model computes.

Two points to remember about scattering: it is both a *source* and a *sink* for the beam — as a source it counts as emission, as a sink as extinction; and it is angle dependent, unlike isotropic thermal emission, which is why the emission coefficient carries an $\hat{\mathbf{n}}$.

Number densities, cross sections, and level populations

The absorption and scattering coefficients are sums of number density times cross section over all species j and all their internal states i:

$$\boxed{\alpha(T,P,\nu) = \sum_j \alpha_j(T,P,\nu) = \sum_j \sum_i n_{ji}(T,P)\,\sigma_{ji}(T,P,\nu)}$$

and similarly $\sigma_s = \sum n\sigma_{\text{scat}}$ for Rayleigh scattering by molecules and for solid particles.

- T and P vary with location, so (T,P) and $\mathbf{x}$ are interchangeable arguments.
- The number densities come from the elemental abundances, atmospheric chemistry, and escape (Chapter 4); the partition among states i from statistical mechanics (LTE, below).
- The inputs are enormous: for water vapor, *hundreds of millions* of lines from transitions among electronic, vibrational, and rotational states (Heng Chapter 5: the "million- to billion-line challenge").

Mass-independent opacity. Many books use $\kappa_m(T,P,\nu) = \sum_j \kappa_j/\rho_j$ [$\mathrm{m^2\,kg^{-1}}$], the extinction per unit mass of species j (Heng's κ). Its virtue: the number density depends strongly on T and P through the ideal gas law, so $\kappa = n\sigma$ varies enormously with height, whereas $\kappa_m = \sigma/m$ varies only through the cross section itself (line broadening, level populations), much more slowly.

Extinction by particles: the extinction efficiency

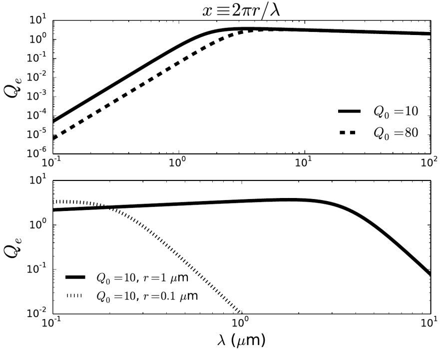

Figure 5: Extinction efficiency of spherical cloud particles versus normalized radius $x = 2\pi r/\lambda$ (top) and versus wavelength for $r = 0.1$ and $1\ \mu$m (bottom). $Q_0 = 10$ and 80 mimic refractory and volatile compositions. Reproduced from Heng (2017), Fig. 2.2.

A spherical particle of radius r has geometric cross section πr^2, but its actual cross section for extinction is generally larger,

$$\boxed{\sigma = Q_e \pi r^2, \qquad Q_e = Q_a + Q_s}$$

with the **extinction efficiency** Q_e, separable into absorption and scattering parts.

- Its functional form and value depend on whether the particle is "small" or "large". "The beauty about physics is that whether something is small or large is not a subjective judgment based on human emotion": for the interaction of radiation with matter, it is a comparison with the wavelength, through the normalized radius

$$x \equiv \frac{2\pi r}{\lambda}.$$

- A particle is small if $x \ll 1$ and large if $x \gg 1$.

Computing Q_e

- To derive the exact functional form of $Q_e(\lambda, r)$ requires a full-blown calculation of absorption and scattering by the particle: Mie theory for homogeneous spheres, given the complex refractive index $n(\lambda) + ik(\lambda)$ of the material, or numerical methods (T-matrix, discrete dipole approximation) for nonspherical particles.
- In practice, we, as atmospheric scientists, look up tables of $Q_e(\lambda, r)$ (and of ω_0 and g_0) and apply them to our models. Optical constants exist for silicates, iron, corundum, water, ammonia, methane, and other candidate condensates.
- Roughly speaking, the extinction efficiency has a somewhat universal form,

$$Q_e \sim \min\left\{\left(\frac{2\pi r}{\lambda}\right)^4, 1\right\},$$

  with the transition from Rayleigh scattering ($x \ll 1$) to scattering by large particles ($x \gg 1$) occurring at a value of the normalized radius that depends on composition.
- With the exception of resonant transitions specific to each material (e.g. the 10 μm silicate feature), the overall shape of Q_e is similar for different members of the same family, e.g. the olivine group.

A fitting formula for Q_e

- A crude but useful **fitting formula** (Heng Eq. 2.25):

$$\boxed{Q_e \simeq \frac{5}{Q_0\, x^{-4} + x^{0.2}}} \qquad x = \frac{2\pi r}{\lambda},$$

  which asymptotes to a roughly constant value for $x \gg 1$ and to Rayleigh scattering ($Q_e \propto x^4$) for $x \ll 1$.
- The dimensionless parameter Q_0 is a proxy for composition: refractory species (e.g. silicates) tend to have $Q_0 \approx 10$, while volatile ones (e.g. water) have $Q_0 \approx 40$–80. A larger Q_0 shifts the Rayleigh-to-geometric transition to larger x.
- The top panel of Heng Fig. 2.2 holds for *all* particle radii: as a function of the normalized radius x, the extinction efficiency is size independent to lowest order. The bottom panel shows the same curve mapped to wavelength for $r = 0.1$ and $1\ \mu$m.
- Given Q_e, a cloud with number density n_c of particles of radius r contributes an extinction coefficient $\kappa_c = n_c Q_e \pi r^2$, to be added to the gas opacity in the radiative transfer; for a size distribution $n_c(r)$ one integrates over r.

Small and large particles

- $x \ll 1$**: Rayleigh regime.** $Q_e \propto x^4$, i.e. $\sigma \propto r^6/\lambda^4$: small particles scatter blue light strongly and are transparent in the infrared. This is the origin of the blue sky, of the Rayleigh slope in transmission spectra, and of the "haze" signature at short wavelengths.
- $x \gg 1$**: geometric regime.** $Q_e \to$ a roughly constant value ($\simeq 2$ for total extinction, the "extinction paradox": geometric blocking plus an equal amount of diffraction); the particle is *grey*, affecting all wavelengths equally.
- A micron-sized particle is radiatively *large* for visible light but *small* in the far-infrared. Generally the influence of a cloud particle on absorption and scattering is felt only at $\lambda \lesssim 2\pi r$, which generalizes straightforwardly to a cloud with a size distribution.
- The particle radius exerts a somewhat larger influence on Q_e than the composition (bottom panel of Heng Fig. 2.2).
- **Consequence for spectra:** a cloud of $\sim 1\,\mu$m particles adds a grey continuum opacity in the optical and near-infrared, flattening the transmission spectrum and muting molecular features, while leaving the mid-infrared unaffected. Sub-micron hazes instead produce a steep, Rayleigh-like slope.

Single-scattering albedo and asymmetry factor

Distinguishing the absorption and scattering cross sections, the **single-scattering albedo** is

$$\boxed{\omega_0 \equiv \frac{\sigma_s}{\sigma_s + \sigma_a}} \qquad \text{(Seager's } \tilde{\omega}\text{, and in coefficient form } \omega_0 = \sigma_s/\kappa\text{),}$$

the fraction of light scattered rather than absorbed in a single interaction: $\omega_0 = 1$ is pure (conservative) scattering, $\omega_0 = 0$ pure absorption.

The second quantity needed to describe a scattering event is the **scattering asymmetry factor** $g_0 = \langle \cos\Theta \rangle$, the phase-function-weighted mean cosine of the scattering angle (Heng Chapter 3):

- $g_0 = 0$: symmetric scattering (isotropic, or Rayleigh, which is symmetric fore–aft but not isotropic); $g_0 > 0$: forward peaked; $g_0 < 0$: backscattering.
- Asymptotically $g_0 \to 0$ for $2\pi r/\lambda \ll 1$ and $g_0 \to 1$ for $2\pi r/\lambda \gg 1$: small particles scatter isotropically, large particles deflect radiation mostly forward. A sub-micron particle scatters infrared light isotropically but forward-scatters X-rays.

Both $\omega_0(\lambda, r)$ and $g_0(\lambda, r)$ are tabulated from Mie calculations for each composition. An atmosphere contains an enormous number of particles of different species; collectively they set the fraction of light reflected at each wavelength, i.e. the albedos of last lecture. Heng Chapter 4 relates A_{B} to ω_0 and g_0 in the two-stream approximation.

Mean opacities: why a plain average is useless

In many treatments of radiative transfer we want a *mean* opacity so that we can work with equations averaged over frequency (grey treatments, Heng Chapter 4). At first thought, one might take the literal average of the opacity over all frequencies. In practice this would not be useful.

- Consider a planet radiating approximately as a blackbody, with exceedingly high opacity at UV frequencies but close to zero opacity everywhere else. A plain average would be dominated by the UV, where there is virtually no planetary flux; a better representation of this atmosphere would simply be zero opacity.
- The point is that we are interested in how much *flux* is blocked by, or allowed through, the atmosphere at different frequencies. Useful mean opacities are therefore weighted by a function of the intensity.
- Because the intensity of the planet atmosphere is unknown — it is the quantity we are trying to solve for — the blackbody intensity is used as the weight.

Planck mean, weighted by B:

$$\boxed{\kappa_{\mathrm{P}}(T,P) = \frac{\int_0^\infty \kappa(T,P,\nu)\,B(T,\nu)\,\mathrm{d}\nu}{\int_0^\infty B(T,\nu)\,\mathrm{d}\nu}}$$

Where the opacity is high, the contribution to the mean is high. The Planck mean is valid in *optically thin* regimes: it gives the total thermal emission of a thin layer correctly, since $\int \alpha_\nu B_\nu\,\mathrm{d}\nu = \kappa_{\mathrm{P}}\,\sigma_{\mathrm{R}} T^4/\pi$.

The Rosseland mean

Rosseland mean, a *harmonic* mean weighted by the temperature derivative of the Planck function:

$$\boxed{\frac{1}{\kappa_{\mathrm{R}}(T,P)} = \frac{\int_0^\infty \kappa(T,P,\nu)^{-1}\,[\partial B(T,\nu)/\partial T]\,\mathrm{d}\nu}{\int_0^\infty [\partial B(T,\nu)/\partial T]\,\mathrm{d}\nu}}$$

- It gives a *higher* weight to frequencies with *small* opacity than to frequencies with large opacity, capturing the physical situation that more radiation travels through the atmosphere at the frequencies where the opacity is smallest — the spectral windows. This is the opposite of the Planck mean.
- It follows from *radiative diffusion* (Chapter 6): deep in an optically thick atmosphere the flux at each frequency is $F_\nu = -\frac{4\pi}{3\kappa_\nu}\frac{\partial B_\nu}{\partial T}\frac{\mathrm{d}T}{\mathrm{d}z}$, and integrating over frequency gives $F = -\frac{16\sigma_{\mathrm{R}}T^3}{3\kappa_{\mathrm{R}}}\frac{\mathrm{d}T}{\mathrm{d}z}$: the total flux is carried where $1/\kappa_\nu$ is large, so the harmonic mean is the right average.
- The weight $\partial B/\partial T$ peaks at slightly shorter wavelengths than B itself ($\lambda T \simeq 2400\,\mu\mathrm{m\,K}$), because the Wien side responds more strongly to a temperature change.

Typically the Planck mean is valid in optically thin regions of the atmosphere and the Rosseland mean in optically thick regimes; neither is right in between, which is why grey models are only qualitative and why Heng Chapter 5 discusses other means and the k-distribution method.

Optical depth, formally, and the plane-parallel convention

From our viewpoint distance is measured in km; for photons the natural distance scale is the optical depth, a measure of transparency. Along a path s,

$$\boxed{\mathrm{d}\tau(s,\nu) = -\kappa(s,\nu)\,\mathrm{d}s}$$

dimensionless; $\tau \ll 1$ optically thin (a photon typically crosses without interacting), $\tau > 1$ optically thick. Because κ depends on frequency and position, so does $\tau(\mathbf{x},\nu)$.

- In **plane-parallel** atmospheres τ is measured *backward* along the outgoing ray: $\tau = 0$ at the top of the atmosphere, increasing downward, because the observer looks down into the atmosphere. Hence the minus sign:

  $$\tau(z,\nu) = -\int_{z_{\max}}^{z} \kappa(z',\nu)\,\mathrm{d}z' = \int_{z}^{z_{\max}} \kappa(z',\nu)\,\mathrm{d}z'.$$

  We write τ_ν with the z dependence implied.
- Later we will see that τ is the dimensionless *e-folding* scale for the attenuation of hot radiation through a cooler gas layer.
- Seager's κ is Heng's α_e; in mass units, $\mathrm{d}\tau = \kappa_m \rho\,\mathrm{d}s$, and with hydrostatic equilibrium $\mathrm{d}\tau = \kappa_m\,\mathrm{d}P/g$.

The mean free path

Consider a slab of area A^2, thickness $\mathrm{d}z$, and number density n, each particle presenting a cross section σ. The probability that an incident photon is stopped in the slab is the ratio of the net area of the targets to the slab area,

$$P(z) = \frac{n\sigma A^2\,\mathrm{d}z}{A^2} = n\sigma\,\mathrm{d}z,$$

which is the same as the fractional drop of the beam intensity that we have already described, $\mathrm{d}I = -\kappa I\,\mathrm{d}z = -n\sigma I\,\mathrm{d}z$. This suggests a reasonable definition of the **mean free path**, the mean distance a photon travels before it is stopped by an interaction with a molecule or other atmospheric particle:

$$\boxed{l = \frac{1}{n\sigma} = \frac{1}{\kappa}, \qquad \mathrm{d}\tau = \frac{\mathrm{d}z}{l}}$$

- The optical depth counts the path length in units of mean free paths; $\tau \sim 1$ means "about one interaction", as Heng's intuitive argument said.
- Typical values: in Earth's lower atmosphere at visible wavelengths (Rayleigh scattering, $\sigma \simeq 5\times10^{-27}\,\mathrm{cm}^2$, $n \simeq 2.5\times10^{19}\,\mathrm{cm}^{-3}$) $l \simeq 80\,\mathrm{km}$, so the clear sky has $\tau \simeq 0.1$ vertically; inside a water cloud l is tens of meters; in the $15\,\mu\mathrm{m}$ CO_2 band core l is meters.

The mean free path, more formally

Start from $\mathrm{d}I = -\kappa I\,\mathrm{d}z$, i.e.

$$\frac{\mathrm{d}I}{I} = -\frac{\mathrm{d}z}{l} \qquad \Longrightarrow \qquad I = I_0\,\mathrm{e}^{-z/l}.$$

The probability that a photon is absorbed between z and $z+\mathrm{d}z$ is the fraction of the initial beam removed there,

$$\mathrm{d}P(z) = \frac{I(z) - I(z+\mathrm{d}z)}{I_0} = \frac{1}{l}\,\mathrm{e}^{-z/l}\,\mathrm{d}z,$$

a properly normalized probability distribution ($\int_0^\infty \mathrm{d}P = 1$). Its expectation value is

$$\boxed{\langle z\rangle = \int_0^\infty z\,\mathrm{d}P(z) = \int_0^\infty \frac{z}{l}\,\mathrm{e}^{-z/l}\,\mathrm{d}z = l}$$

The mean distance traveled before an interaction is indeed $l = 1/\kappa$, which justifies the name.

- In optical-depth units the distribution is $\mathrm{d}P = \mathrm{e}^{-\tau}\mathrm{d}\tau$, with mean $\langle\tau\rangle = 1$ and standard deviation 1: interaction distances are broadly distributed, and a sizable fraction of photons travel $\tau > 2$ or 3 without interacting ($\mathrm{e}^{-2} = 14\%$, $\mathrm{e}^{-3} = 5\%$).
- Monte Carlo radiative transfer samples exactly this distribution: $\tau = -\ln\xi$ with ξ a uniform random number in $(0,1]$, then the photon is moved to the point where the accumulated optical depth along its direction equals τ.

Local thermodynamic equilibrium (LTE)

Complete thermodynamic equilibrium requires thermal, chemical, and mechanical equilibrium — impossible across an atmosphere with huge gradients of T and P and an open boundary at the top. **Local** thermodynamic equilibrium (LTE) applies in a local region where the gradients are small compared to the photon mean free path.

- In LTE we assume all the conditions of thermodynamic equilibrium hold for the *matter*, but we let the *radiation field* depart from a blackbody. The radiation field in a planetary atmosphere is very different from B; it must be computed from the transfer equation.
- Why it matters: the coefficients κ and ε are sums of number densities times cross sections over all **level populations** (the numbers of atoms or molecules in each quantum state). We need (1) the overall number densities and (2) the fraction in each state.
- The heart of the problem: the level populations are determined by the radiation field (photoexcitation, stimulated emission, photoionization) — the very quantity we are solving for. A full simultaneous solution of the radiation field and the populations is possible in principle but onerous in practice.
- **LTE decouples them:** all properties of the matter depend only on the local kinetic temperature and one other variable (ρ or P). Level populations follow the Boltzmann distribution, ionization the Saha equation, molecular abundances chemical equilibrium (Chapter 8 and Heng Chapter 7).

When is LTE valid?

LTE implies a strict connection between the matter and the local temperature; it is valid where that connection is physically realized.

- **High densities: collisions dominate radiative processes.** An electron in an atom or molecule is excited to a higher level by absorbing a photon. Later, because of a collision, it cascades downward, releasing the energy as kinetic energy of the colliding particles; after further elastic collisions this energy ends up in the thermal pool of the gas. In this case of collisional deexcitation the photon is said to be *destroyed* and converted into kinetic energy. In this way collisions couple the photons (the radiation field) to the matter temperature via the kinetic energy of the gas. (As long as the colliding particles have a Maxwellian velocity distribution, the level populations have their Boltzmann equilibrium values.)
- **A blackbody radiation field.** If the mean radiation field is a blackbody, $J(\mathbf{x}, \nu) = B(\mathbf{x}, \nu)$, the radiative rate equations that control the level populations also give an equilibrium (Boltzmann) distribution (Section 8.2.4). This second situation is not really relevant for planetary atmospheres, and from the first example it is clearly not a *necessary* condition for LTE.
- The criterion is a comparison of rates: LTE holds for a transition when the collisional deexcitation rate exceeds the radiative one, $n_{\mathrm{coll}}\, q_{ul} \gg A_{ul}$, i.e. above a *critical density* that depends on the transition (Chapter 8).

LTE in Earth's atmosphere and in exoplanets

- **Earth:** LTE is actually valid from the ground up to about 60 km, largely because collisional rates between molecules dominate radiative rates, driving a Boltzmann distribution. We know from spectral measurements at the top of the atmosphere and at different altitudes that Earth's radiation field is *not* a blackbody — LTE and blackbody radiation are different things. Above ~ 60 km (mesosphere, thermosphere) the CO_2 15 μm band and other bands are in non-LTE: collisions are too rare to keep the vibrational levels thermalized.
- **Exoplanets:** LTE is commonly used in exoplanet atmosphere calculations, and it is a good approximation for the pressures probed by emission spectra ($\gtrsim 1$ mbar). But there are regions where it is not valid: the very upper layers where radiation escapes freely through the open boundary. There the radiation field is very different from a blackbody *and* collisions do not dominate over radiative transition rates, causing a departure from LTE. Examples: CH_4 fluorescence at 3.3 μm in giant planets, the escaping upper atmospheres probed by UV transits, and the thermospheres heated by stellar XUV.
- **Validation:** for a given situation one could compare radiative and collisional rates, but to accurately validate LTE a non-LTE calculation itself may be necessary. In practice one checks whether the pressures probed by the observation are above the critical densities of the relevant transitions.

LTE and Kirchhoff's law

Kirchhoff's law of radiation states that in thermal equilibrium the emissivity of a body equals its absorptivity:

$$\boxed{\varepsilon_{\text{therm}}(\mathbf{x},\nu) = \alpha(\mathbf{x},\nu)\,B(\mathbf{x},\nu)}$$

- The thermal emission at $\mathbf{x}$ is blackbody radiation weighted by the absorptive properties of the gas: any time a photon is absorbed by gas in LTE, the energy is reemitted as radiation, with amount and wavelength set by the Planck function times the absorption coefficient. Gas that absorbs strongly at a frequency also emits strongly there — spectral lines appear in absorption or emission depending on the temperature gradient, not on the opacity alone.
- Why it must hold (exercise 5.1): inside a blackbody enclosure at temperature T the radiation field is B and is unchanged by the presence of any body at the same T; the transfer equation $\mathrm{d}I/\mathrm{d}s = -\kappa I + \varepsilon$ with $I = B$ constant then demands $\varepsilon = \kappa B$ for every frequency and direction.
- Strictly valid only in complete thermodynamic equilibrium. In planetary atmospheres we justify it for small, localized regions where the gradients of the thermodynamic properties are small compared to a photon mean free path.
- Note what Kirchhoff's law does *not* say: it does not make $I = B$. The emissivity is tied to the local temperature; the intensity is whatever the transfer equation produces.

The source function

The **source function** (sometimes called the contribution function) is the ratio of the emission coefficient to the extinction coefficient:

$$\boxed{S(\mathbf{x}, \hat{\mathbf{n}}, \nu) = \frac{\varepsilon(\mathbf{x}, \hat{\mathbf{n}}, \nu)}{\kappa(\mathbf{x}, \nu)}}$$

It has the same dimensions as the intensity, it simplifies the transfer equation and its solution, and it provides physical insight: it is the intensity that the beam *relaxes toward* as it travels through the medium (a beam brighter than S is attenuated, a beam fainter than S is filled in).

Expand the source function into a thermal-emission component and a scattering component. With Kirchhoff's law $\varepsilon_{\text{therm}} = \alpha B$ and isotropic coherent scattering $\varepsilon_{\text{scat}} = \sigma_s J$, under conditions of LTE,

$$S(\mathbf{x}, \nu) = \frac{\alpha(\mathbf{x}, \nu) B(\mathbf{x}, \nu) + \sigma_s(\mathbf{x}, \nu) J(\mathbf{x}, \nu)}{\alpha(\mathbf{x}, \nu) + \sigma_s(\mathbf{x}, \nu)} \quad \Longleftrightarrow \quad \boxed{S = (1 - \omega_0) B + \omega_0 J}$$

with the single-scattering albedo $\omega_0 = \sigma_s / (\alpha + \sigma_s)$: a weighted mean of the Planck function and the mean intensity.

The source function: limits and consequences

- **Pure absorption** ($\omega_0 = 0$): the strict LTE value $S = B$, from Kirchhoff's law with no scattering. The beam relaxes toward the local Planck function; the emergent spectrum then reflects the temperature structure directly.
- **Pure scattering** ($\omega_0 = 1$): $S = J$, the mean intensity. The beam relaxes toward the local radiation field itself; the medium neither creates nor destroys photons, it only redistributes their directions.
- The scattering term makes S depend on the radiation field: the transfer equation becomes an *integro-differential* equation for I. Even the expression with $\omega_0 \neq 0$ "allows a slight deviation from LTE" in Seager's words: the level populations are still thermal, but the source function is not B.
- In general, while $I = B$ for blackbody radiation, LTE radiation is *not* necessarily blackbody radiation: in LTE, I is not required to equal B. LTE constrains the *matter* (Kirchhoff, Boltzmann, Saha); the *radiation* is whatever the transfer equation produces.
- **Convention check.** Heng writes $S = \omega_0 J/4\pi + (1-\omega_0)B$ because his $J \equiv \oint I\,\mathrm{d}\Omega$ is 4π times Seager's mean intensity $J = \frac{1}{4\pi}\oint I\,\mathrm{d}\Omega$. Same physics; check which J a paper uses before comparing formulas.

The radiative transfer equation along a ray

Radiative transfer describes how a beam changes as it travels a distance s through a volume of gas: losses κI and gains ε. For a static atmosphere,

$$\boxed{\frac{\mathrm{d}I(\mathbf{x},\hat{\mathbf{n}},\nu)}{\mathrm{d}s} = -\kappa(\mathbf{x},\nu)\,I(\mathbf{x},\hat{\mathbf{n}},\nu) + \varepsilon(\mathbf{x},\hat{\mathbf{n}},\nu)}$$

Time-dependent derivation. In an inertial frame, follow the energy in a fixed volume element of length ds and cross section dA during dt; for moving material κ and ε acquire an angle dependence through the Doppler shift. Balancing losses and gains,

$$[I(\mathbf{x}+\Delta\mathbf{x},\hat{\mathbf{n}},\nu,t+\Delta t) - I(\mathbf{x},\hat{\mathbf{n}},\nu,t)]\,\mathrm{d}A\,\mathrm{d}\Omega\,\mathrm{d}\nu\,\mathrm{d}t = [-\kappa I + \varepsilon]\,\mathrm{d}s\,\mathrm{d}A\,\mathrm{d}\Omega\,\mathrm{d}\nu\,\mathrm{d}t,$$

and with $\Delta t = \Delta s/c$ the left side is $\left(\frac{1}{c}\partial_t I + \partial_s I\right)\mathrm{d}s$, giving

$$\left(\frac{1}{c}\frac{\partial}{\partial t} + \frac{\partial}{\partial s}\right) I(\mathbf{x},\hat{\mathbf{n}},\nu,t) = -\kappa(\mathbf{x},\hat{\mathbf{n}},\nu,t)\,I + \varepsilon(\mathbf{x},\hat{\mathbf{n}},\nu,t).$$

- The ∂_t/c term matters only when conditions change on light-crossing times; for atmospheres we drop it (static case), together with the $\hat{\mathbf{n}}$ dependence of κ, while remembering that radiation is an energy *flow* per unit time.
- In vacuum ($\kappa = \varepsilon = 0$) the equation says I is constant along the ray, as claimed last lecture.

Heng's shortest derivation, in optical depth

Since τ is the correct measure of transparency, work directly with it. Let the intensity impinge on the surface of an object at angle θ from the normal, $\mu \equiv \cos\theta$, with the convention that reflection (upward travel) has $\mu > 0$ and penetration into the object $\mu < 0$.

- Loss of intensity by extinction across the object's optical depth $\Delta\tau$ (measured along the normal): $I\,\Delta\tau/\mu$ (the slant path is $\Delta\tau/|\mu|$ long).
- Compensated by the object's own emission: $-S\,\Delta\tau/\mu$.

$$\Delta I = \frac{\Delta\tau}{\mu}(I - S) \quad \xrightarrow{\Delta\tau\to 0} \quad \boxed{\mu\,\frac{\partial I}{\partial\tau} = I - S}$$

- "It looks deceivingly simple, because all of the complexity has been hidden within the source function."
- The same equation follows from Seager's form by substituting $\mathrm{d}\tau = -\kappa\,\mathrm{d}z$, $\mathrm{d}z = \mu\,\mathrm{d}s$, and $S = \varepsilon/\kappa$ (next slides). The sign convention (τ increasing downward, $\mu > 0$ upward) is the standard one for atmospheres.
- Historical note: the equation was first derived for stars by astrophysicists (Schuster, Schwarzschild, Milne, Chandrasekhar) and later adopted by the atmospheric and climate sciences.

The plane-parallel approximation

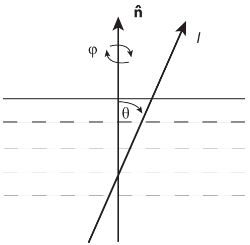

Figure 6: A 1D plane-parallel atmosphere: stratified layers, vertical coordinate z, ray at angle θ from the normal $\hat{\mathbf{n}}$, azimuth ϕ. Reproduced from Seager (2010), Fig. 5.2.

The plane-parallel atmosphere is a good framework in which to study the radiative transfer equation.

- For a 1D planar atmosphere, the atmosphere is defined by a stratified plane with each layer having homogeneous properties such as T, P, and ρ. This 1D plane represents one location on the surface of the planet.
- **When is it valid?** If the radial depth of the atmosphere is much smaller than the planetary radius. The photosphere spans a few scale heights: $H/R_p \sim 0.01$ for hot Jupiters, 10^{-3} for Earth, so the curvature of the layers is negligible for rays that are not nearly horizontal.
- The plane-parallel definition assumes **axial symmetry** about the vertical: there is no preferred horizontal direction in the medium itself (the illumination may still come from a specific direction).
- The vertical coordinate z increases upward; the optical depth τ increases downward from $\tau = 0$ at the top.

The plane-parallel transfer equation

Considering the axial symmetry and adopting z as the (1D) vertical coordinate,

$$\frac{\partial}{\partial x} = \frac{\partial}{\partial y} = 0 \qquad \text{and} \qquad \frac{\mathrm{d}z}{\mathrm{d}s} = \cos\theta \equiv \mu,$$

where θ is the angle between the surface normal and the beam of intensity. Substituting into the equation along a ray,

$$\boxed{\mu \frac{\mathrm{d}I(z,\hat{\mathbf{n}},\nu)}{\mathrm{d}z} = -\kappa(z,\nu)\, I(z,\hat{\mathbf{n}},\nu) + \varepsilon(z,\hat{\mathbf{n}},\nu)}$$

If we further consider axial symmetry of the radiation field, i.e. no ϕ dependence,

$$\mu \frac{\mathrm{d}I(z,\mu,\nu)}{\mathrm{d}z} = -\kappa(z,\nu)\, I(z,\mu,\nu) + \varepsilon(z,\mu,\nu),$$

and a single angle variable $\mu \in [-1,1]$ describes the direction: $\mu > 0$ upward, $\mu < 0$ downward, $\mu = 0$ horizontal.

- The factor μ is the whole geometric content of the approximation: a slant ray traverses $\mathrm{d}s = \mathrm{d}z/\mu$ of path for each $\mathrm{d}z$ of height, so it sees $1/\mu$ times more optical depth.
- Horizontal rays ($\mu \to 0$) travel infinitely far in a plane-parallel atmosphere; in reality they leave the curved atmosphere after a path $\sim \sqrt{2\pi R_p H}$, the transit chord.

The scattering integral and the limits of plane-parallel geometry

As written, the plane-parallel equation is an *integro-differential* equation, because $\varepsilon = \varepsilon_{\text{scat}} + \varepsilon_{\text{therm}}$ and the scattering term contains angle integrals of I:

$$\varepsilon_{\text{scat}}(z,\mu,\nu) = \sigma_s(z,\nu)\,\frac{1}{4\pi}\int_0^{2\pi}\int_{-1}^{1} P(\mu,\phi;\mu',\phi')\,I(z,\mu',\nu)\,\mathrm{d}\mu'\,\mathrm{d}\phi'.$$

Even with azimuthal symmetry of the medium, the phase function couples every direction to every other direction, so the equation for $I(\mu)$ at one angle involves I at all angles.

Where the plane-parallel picture fails.

- **Transmission at the terminator:** the starlight travels along a chord, and the geometry is inherently spherical; the slant path length $\sqrt{2\pi R_p H}$ comes from the curvature.
- **Grazing incidence near the limb** of the dayside, where the stellar zenith angle approaches 90° and $1/\mu_0$ diverges; spherical (or pseudo-spherical) corrections are needed.
- **Extended atmospheres** inflated by escape or by very low gravity, where H/R_p is no longer small.
- **Horizontal inhomogeneity:** day–night contrasts, patchy clouds, and the eclipse-mapping problem require 3D or at least a collection of 1D columns.

Moments and the equation in optical depth

In plane-parallel geometry with azimuthal symmetry the mean intensity and flux of Chapter 2 become

$$J(z,\nu) = \frac{1}{4\pi}\int_0^{2\pi}\int_{-1}^{1} I\,\mathrm{d}\mu\,\mathrm{d}\phi = \frac{1}{2}\int_{-1}^{1} I(z,\mu,\nu)\,\mathrm{d}\mu, \qquad F(z,\nu) = \int_0^{2\pi}\int_{-1}^{1} \mu I\,\mathrm{d}\mu\,\mathrm{d}\phi = 2\pi\int_{-1}^{1}\mu\, I(z,\mu,\nu)\,\mathrm{d}\mu.$$

(The stellar-atmosphere literature uses $H = F/4\pi$ because its form parallels J and K.)

Now change to the optical-depth scale ($\mathrm{d}\tau_\nu = -\kappa\,\mathrm{d}z$) and use the source function:

$$\boxed{\mu\frac{\mathrm{d}I(\tau_\nu,\mu,\nu)}{\mathrm{d}\tau_\nu} = I(\tau_\nu,\mu,\nu) - S(\tau_\nu,\mu,\nu)}$$

the same form Heng obtained directly. Note the roles of the two hemispheres: **outgoing** rays have $\mu > 0$, **incoming** rays $\mu < 0$.

Boundary conditions.

- Top ($\tau_\nu = 0$): the stellar radiation incident on the atmosphere, $I(0,\mu,\nu) = I_*(0,\mu_0,\nu)$ for $-1 \le \mu \le 0$, a plane-parallel beam at the single angle μ_0. For the nightside, for infrared calculations where the planet's own emission dominates, or for unirradiated planets, $I_* = 0$.
- Bottom ($\tau_{\mathrm{max},\nu}$): the intensity from the planet's interior, $I(\tau_{\mathrm{max},\nu},\mu,\nu) = I_{\mathrm{int}}$ for $0 \le \mu \le 1$, matched to the interior energy flux (or, for a terrestrial planet, the surface emission and reflection).

The formal solution

Multiply the equation by the integrating factor $\mathrm{e}^{-\tau_\nu/\mu}$:

$$\frac{\mathrm{d}I}{\mathrm{d}\tau_\nu}\mathrm{e}^{-\tau_\nu/\mu} - \frac{I}{\mu}\mathrm{e}^{-\tau_\nu/\mu} = \frac{\mathrm{d}}{\mathrm{d}\tau_\nu}\left(I\mathrm{e}^{-\tau_\nu/\mu}\right) = -\frac{S}{\mu}\mathrm{e}^{-\tau_\nu/\mu}.$$

Integrating from an initial optical depth $\tau_{\nu,i}$ to a final one $\tau_{\nu,f}$,

$$\boxed{I(\tau_{\nu,f},\mu,\nu) = I(\tau_{\nu,i},\mu,\nu)\,\mathrm{e}^{-(\tau_{\nu,i}-\tau_{\nu,f})/\mu} - \frac{1}{\mu}\int_{\tau_{\nu,i}}^{\tau_{\nu,f}} S(\tau_\nu,\mu,\nu)\,\mathrm{e}^{-(\tau_\nu-\tau_{\nu,f})/\mu}\,\mathrm{d}\tau_\nu}$$

- **First term:** the initial intensity diminished by the exponential attenuation of absorption over the optical path $(\tau_i - \tau_f)/\mu$.
- **Second term:** the emission from the atmosphere, an exponentially weighted average of the source function along the beam up to the location of interest.
- For an outgoing ray ($\mu > 0$) from deep inside ($\tau_i > \tau_f$) the integral runs downward in τ and the second term is positive: each layer contributes $S\,\mathrm{d}\tau/\mu$, attenuated by $\mathrm{e}^{-(\tau-\tau_f)/\mu}$ on its way up. For an incoming ray ($\mu < 0$) the same formula, read from the top down, describes the penetration of starlight with the atmosphere's emission added.

It is called the "formal" solution because S must be known; in general it is not.

The emergent intensity

Our goal is the planet's spectrum: the emergent flux at the top of the atmosphere, the measurable quantity. For a *semi-infinite* atmosphere, integrate from deep inside, $\tau_{\nu,i} = \infty$, to the top, $\tau_{\nu,f} = 0$, using $\lim_{\tau_\nu \to \infty} I e^{-\tau_\nu/\mu} = 0$ (the deep layers are hidden):

$$\boxed{I(0,\mu,\nu) = \frac{1}{\mu}\int_0^\infty S(\tau_\nu,\mu,\nu)\,\mathrm{e}^{-\tau_\nu/\mu}\,\mathrm{d}\tau_\nu} \qquad (0 < \mu \leq 1).$$

The emergent intensity is the amount of intensity from each altitude that reaches the surface along a path at angle θ to the line of sight.

- The weight $\mathrm{e}^{-\tau/\mu}/\mu$ is normalized ($\int_0^\infty \mathrm{e}^{-\tau/\mu}\mathrm{d}\tau/\mu = 1$) and peaks at the top, with mean $\langle\tau_\nu\rangle = \mu$: the emergent intensity is a weighted sample of S over $\tau_\nu \lesssim \mu$, the "photosphere" at that frequency and angle. Rays emerging at grazing angles ($\mu \to 0$) sample shallower, usually cooler layers — this is **limb darkening**.

Eddington–Barbier and the formation of bands

- For a source function linear in optical depth, $S = a + b\tau$, the emergent-intensity integral gives *exactly*

$$I(0,\mu) = \frac{1}{\mu}\int_0^\infty (a+b\tau)\,\mathrm{e}^{-\tau/\mu}\,\mathrm{d}\tau = a + b\mu = S(\tau=\mu),$$

  the **Eddington–Barbier relation**: the emergent intensity equals the source function at optical depth μ along the ray. For a general S it remains a good approximation.
- **Limb darkening.** Looking at the disk center ($\mu = 1$) we see $S(\tau = 1)$; near the limb ($\mu \to 0$) we see $S(\tau \to 0)$. If T decreases outward, the limb is darker; a stellar limb darkening curve therefore measures $\mathrm{d}S/\mathrm{d}\tau$, and modeling it is essential for transit radii (Lecture 1).
- **Absorption bands.** In a strong molecular band κ_ν is large, $\tau_\nu = \mu$ is reached high up where T is low, and $I(0,\mu,\nu) = B(T_{\mathrm{high}})$ is small: the band appears in *absorption*. In a spectral window we see deeper, warmer layers and more flux.
- **Emission bands.** If T *increases* with height (a stratospheric inversion, as for ozone on Earth or TiO/VO on some hot Jupiters), the same bands appear in *emission*. The sign of the features in an emission spectrum thus diagnoses the sign of the temperature gradient.
- This is the whole basis of remote sensing of temperature structure: each frequency samples S, hence T, at its own $\tau_\nu \sim \mu$, i.e. its own pressure (Chapter 6 and Chapter 9).

The emergent flux

The emergent surface flux follows from the flux definition of Chapter 2 applied to the outgoing hemisphere, $F = 2\pi \int_0^1 \mu\, I(0,\mu)\,\mathrm{d}\mu$ (no incoming radiation at the top in the thermal infrared):

$$\boxed{F(0,\nu) = 2\pi \int_0^1 \int_0^\infty S(\tau_\nu,\mu,\nu)\,\mathrm{e}^{-\tau_\nu/\mu}\,\mathrm{d}\tau_\nu\,\mathrm{d}\mu}$$

- If the source function $S(\tau,\mu,\nu)$ is *known*, the emergent intensity and the emergent flux — the planet's spectrum — can be computed directly from these two integrals. The μ integral of $\mathrm{e}^{-\tau/\mu}$ produces the exponential integral $E_2(\tau)$, so $F(0,\nu) = 2\pi \int_0^\infty S(\tau_\nu)E_2(\tau_\nu)\,\mathrm{d}\tau_\nu$ for an isotropic S.
- For $S = B(T(\tau))$ (pure absorption, LTE) this is the standard *forward model* of an emission spectrum: choose $T(\tau_\nu)$, i.e. a temperature–pressure profile and an opacity function, and integrate.
- In many cases, however, this straightforward solution is not possible. The main complication is that the source function itself depends on $I(\tau,\mu,\nu)$, the quantity we are trying to solve for (next slide).

Why the problem is hard

- **The source function depends on I through scattering.** $S = (1-\omega_0)B + \omega_0 J$, or with the angle-dependent scattering emissivity

$$\varepsilon_{\rm scat}(\tau,\mu,\nu) = \sigma_s(\tau,\nu)\,\frac{1}{4\pi}\int_0^{2\pi}\int_{-1}^{1} P(\mu,\phi;\mu',\phi')\,I(\tau,\mu',\nu)\,\mathrm{d}\mu'\,\mathrm{d}\phi',$$

  and J is an integral of the unknown I. Physically, scattering *decouples* the intensity from the local conditions: photons may scatter through large distances in the atmosphere without interacting with the thermal pool of the gas via absorption and thermal reemission. A numerical solution of an integro-differential equation is required.
- **Even B depends on I** through the temperature structure: radiative equilibrium ties $T(\tau)$ to the absorbed radiation (next section), so S is unknown until the whole problem, including the energy balance of every layer, is solved.

The boundary conditions: a two-point problem

- To further investigate the hidden complication, return to the boundary conditions. To solve for the intensity from this *first-order* ordinary differential equation we require two full boundary conditions, that is, an upper and a lower boundary condition on the *full* range $-1 \leq \mu \leq 1$.
- Yet the information we have is the stellar radiation incident on the planet and traveling *downward*, $-1 \leq \mu < 0$, at the top, and an estimate of the interior energy that we may convert to an *outward-going* intensity, $0 < \mu \leq 1$, at the bottom.
- Because the boundary conditions are not fully specified — half of the rays are known at each boundary — iterative techniques or a different formulation of the transfer equation are needed to solve for I:
  - Λ iteration: guess S, compute I by the formal solution, update J and S, repeat (slow when $\omega_0 \to 1$; accelerated versions exist);
  - Feautrier's method: combine $I(\mu)$ and $I(-\mu)$ into symmetric and antisymmetric parts, giving a second-order equation with one condition at each boundary;
  - two-stream and discrete-ordinate methods (Heng Chapter 3), moment methods with a closure (Eddington approximation, Chapter 6), and Monte Carlo simulations.
- The same structure appears in every transport problem with scattering: the solution at a point depends on the solution everywhere else.

Beer's law

The simplest solution discards the source function ($S = 0$): pure extinction of an external beam.

$$\int_{I_0}^{I} \frac{\mathrm{d}I}{I} = \int_0^{\tau} \frac{\mathrm{d}\tau'}{\mu} \quad \Longrightarrow \quad \boxed{I = I_0\,\mathrm{e}^{-\tau/|\mu|}}$$

(Heng writes $I_0 \mathrm{e}^{\tau/\mu}$ with his $\mu < 0$ for penetrating rays; the intensity decreases either way.) This is **Beer's law** (or the Beer–Lambert–Bouguer law): the intensity is diminished exponentially, $\tau = 1$ corresponding to one e-folding, i.e. 63% of the beam removed ($1/\mathrm{e} \simeq 37\%$ transmitted) when unit optical depth is traversed.

- The slant factor $1/|\mu|$: a beam entering at 60° from the normal traverses twice the optical depth of a normal beam. This is why the Sun is redder at sunset and why the transit chord is so opaque.
- Beer's law describes very well how *starlight* is absorbed on entering an exoplanet atmosphere, and it underlies transmission spectroscopy: the in-transit flux is the stellar flux times $\mathrm{e}^{-\tau_{\mathrm{chord}}(\nu)}$ integrated over the annulus.
- It is inadequate for the planet's *thermal emission*, which has both incoming and outgoing components and a source function that cannot be neglected.
- Heng Chapter 4 generalizes Beer's law to include non-isotropic scattering of the stellar beam.

Direct solution with pure absorption: a pair of layers

Beyond Beer's law, the equation can be solved directly for I only in the limit of *pure absorption*. Include coherent isotropic scattering, $S = (1-\omega_0)B + \omega_0 J$, and integrate between two points 1 and 2 ($\tau_2 > \tau_1$):

$$I_2 \mathrm{e}^{-\tau_2/\mu} - I_1 \mathrm{e}^{-\tau_1/\mu} = -\frac{1}{\mu}\int_{\tau_1}^{\tau_2} \left[\omega_0 J + (1-\omega_0)B\right] \mathrm{e}^{-\tau/\mu}\, \mathrm{d}\tau.$$

One cannot proceed unless $\omega_0 = 0$, because $J(\tau)$ is not known a priori. ("Coherent" means the wavelength is unchanged; spectral lines are better described by complete redistribution, but for bins wide compared to line widths coherent scattering is a decent approximation.)

Isothermal absorbing layer. Take $\omega_0 = 0$ and B, I_1, I_2 constant; label upward and downward radiation by $\uparrow$, $\downarrow$; integrate over the hemisphere and write fluxes as $F = \pi I$. Then

$$\boxed{F_{\uparrow 1} = F_{\uparrow 2}\, \mathscr{T} + \pi B\,(1-\mathscr{T}), \qquad F_{\downarrow 2} = F_{\downarrow 1}\, \mathscr{T} + \pi B\,(1-\mathscr{T})}$$

with the **transmission function**

$$\mathscr{T} \equiv 2\int_0^1 \mu\, \mathrm{e}^{-\Delta\tau/\mu}\, \mathrm{d}\mu = 2E_3(\Delta\tau) = (1-\Delta\tau)\mathrm{e}^{-\Delta\tau} + \Delta\tau^2 E_1(\Delta\tau), \qquad \Delta\tau \equiv \tau_2 - \tau_1,$$

where $E_n(x) = \int_1^\infty t^{-n} \mathrm{e}^{-xt}\, \mathrm{d}t$ are the exponential integrals (Heng, problem 2.8.5).

The transmission function

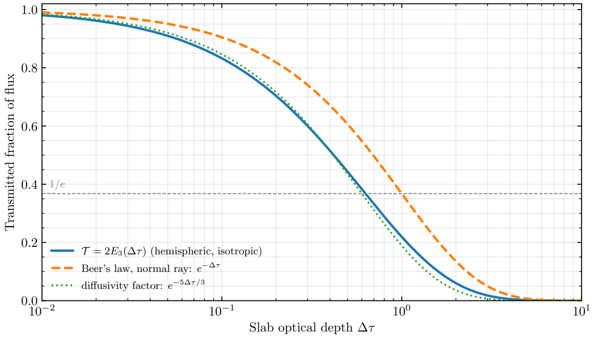

Figure 7: Hemispheric transmission $\mathscr{T} = 2E_3(\Delta\tau)$ of an isothermal slab compared with Beer's law for a normal ray and with the diffusivity-factor approximation $e^{-5\Delta\tau/3}$.

- $\mathscr{T}$ has a convenient range between 0 and 1. When $\mathscr{T} = 1$ no radiation is absorbed and no blackbody emission is produced: the incident flux passes right through the pair of layers, a completely transparent medium. When $\mathscr{T} = 0$ nothing passes through and the pair of layers radiates purely as a blackbody, $F = \pi B$.
- Because the flux is an average over all angles, and slant rays see a larger optical depth $\Delta\tau/\mu$, $\mathscr{T}$ falls faster than the normal-ray Beer's law $e^{-\Delta\tau}$: the slab reaches $1/e$ transmission at $\Delta\tau \simeq 0.6$, not at 1.

The transmission function: approximations and use

- **Diffusivity factor.** The classic approximation replaces the angle integral by a single ray at an effective angle, $\mathscr{T} \simeq \mathrm{e}^{-D\Delta\tau}$ with $D = 5/3$ ($\mu_{\mathrm{eff}} = 3/5$; Elsasser 1942). As the figure shows it is accurate to a few percent for all $\Delta\tau$, which is why single-angle (two-stream) methods work so well for thermal radiation.
- **Small $\Delta\tau$:** $\mathscr{T} \simeq 1 - 2\Delta\tau$ (not $1 - \Delta\tau$): a thin slab absorbs twice as much hemispheric flux as a normal ray would suggest, because the mean slant factor $\langle 1/\mu \rangle$ over the flux-weighted hemisphere is 2.
- **Large $\Delta\tau$:** $\mathscr{T} \simeq \mathrm{e}^{-\Delta\tau}(1 - \Delta\tau) + \Delta\tau^2 E_1 \to 2\mathrm{e}^{-\Delta\tau}/\Delta\tau$ asymptotically, dominated by the normal rays, the only ones that get through.
- **Pairs of layers.** Writing the solutions in this form describes the transfer of radiation from point 1 to point 2 and vice versa, using either location as the starting point or boundary condition. Formulating the problem as radiative transfer between *pairs of layers* is the basis of the two-stream treatment (Heng Chapter 3), and stacking many such layers is how a numerical longwave code computes the emergent flux for an arbitrary $T(\tau)$ profile.

Non-isothermal layers: a linear source function

The Planck function generally varies with optical depth. A direct solution still exists if B is expanded to first order about the layer boundaries (Heng, problem 2.8.5b):

$$B = \begin{cases} B_1 + B'(\tau - \tau_1), & \downarrow \text{ direction}, \\ B_2 + B'(\tau - \tau_2), & \uparrow \text{ direction}, \end{cases} \qquad B' \equiv \frac{\partial B}{\partial \tau} = \frac{B_2 - B_1}{\tau_2 - \tau_1},$$

which reproduces B_1 at τ_1 and B_2 at τ_2. The hemispherically integrated solutions become

$$F_{\uparrow 1} = F_{\uparrow 2}\mathscr{T} + \pi B_2(1-\mathscr{T}) + \pi B'\left[\tfrac{2}{3}\left(1-\mathrm{e}^{-\Delta\tau}\right) - \Delta\tau\left(1-\tfrac{\mathscr{T}}{3}\right)\right],$$

$$F_{\downarrow 2} = F_{\downarrow 1}\mathscr{T} + \pi B_1(1-\mathscr{T}) - \pi B'\left[\tfrac{2}{3}\left(1-\mathrm{e}^{-\Delta\tau}\right) - \Delta\tau\left(1-\tfrac{\mathscr{T}}{3}\right)\right].$$

- The extra term is the correction from the temperature gradient within the layer; it vanishes for $B' = 0$ and for $\Delta\tau \to 0$.

Why the linear source function matters

- The “linear-in-τ” source function is the standard building block of numerical longwave radiative transfer, for example in the widely used two-stream scheme of Toon et al. (1989) and in most exoplanet forward models: within each layer B is interpolated linearly in τ, and the layer solutions are chained together from the bottom boundary to the top.
- Integrated over a semi-infinite atmosphere, it gives the **Eddington–Barbier relation for the flux**: for $S = a + b\tau$, $F(0) = 2\pi \int_0^1 \mu(a + b\mu)\,\mathrm{d}\mu = \pi(a + \frac{2}{3}b) = \pi S(\tau = \frac{2}{3})$. The emergent flux equals π times the source function at optical depth $2/3$ — the origin of the rule of thumb that the photosphere sits at $\tau \simeq 2/3$, and of the appearance of $2/3$ in grey-atmosphere temperature profiles.
- The same relation for the intensity gives $I(0,\mu) = S(\tau = \mu)$ (previous section), so the disk-center intensity samples $\tau = 1$ and the limb samples $\tau \to 0$: limb darkening measures $\mathrm{d}S/\mathrm{d}\tau$, i.e. the temperature gradient.
- With many thin layers, each isothermal or linear, the emergent flux of an arbitrary $T(\tau)$ profile can be built up layer by layer — exactly the forward model used in atmospheric retrievals (below), where speed matters because it is evaluated millions of times.

A practical checklist for radiative transfer calculations

What is needed to actually compute a spectrum (Heng, Section 2.5):

1. **Choose the form of the transfer equation** and the approximations (plane-parallel? scattering? two-stream or multi-stream?), check that the assumptions are plausible, and select the numerical method. Solutions are called *forward models*: they take abundances and opacities and produce a synthetic spectrum.
2. **Set up a model atmosphere** with a finite number of layers, each with its own T, P, and column mass, and compute the fluxes passing through each layer. The isothermal-layer solution above is a forward model in the pure-absorption limit; its inputs are the opacities, from which the optical depths, transmission functions, and fluxes of each layer follow.
3. **Decide the chemical composition** (Heng Chapter 7) and obtain the *opacities* of the atoms, molecules, and aerosols, typically from precomputed tables (Heng Chapter 5). Combining them, weighted by abundance, over T, P, and λ gives the opacity function of the atmosphere.
4. **Iterate to convergence**: opacities, optical depths, and temperatures are interdependent. Convergence is reached when the temperature of each layer stops changing (to within a numerical threshold).

The converged solution yields the flux at the top of the atmosphere (the synthetic spectrum) and the temperatures of all layers (the temperature–pressure profile).

Radiative equilibrium: local energy conservation

How are the temperatures and fluxes of the layers related? In the absence of atmospheric dynamics one hopes to solve for **radiative equilibrium**, which occurs when the net heat entering and leaving each layer vanishes. It is a statement of *local* energy conservation and originates from the first law of thermodynamics (Heng Chapter 9),

$$\rho c_P \frac{DT}{Dt} = \frac{DP}{Dt} + \rho Q,$$

with ρ the mass density, c_P the specific heat at constant pressure, and Q the heating term. Focusing on radiative heating (ignoring convection and conduction),

$$\rho Q = -\nabla \cdot \mathbf{F}_-,$$

where $\mathbf{F}_-$ is the wavelength-integrated *net* flux: the difference between the flux entering and exiting a layer, hence the minus subscript.

- Geometrically this makes sense: when radiation emanates from the layer, the divergence of the net flux is positive, which leads to cooling ($Q < 0$). The opposite happens when radiation enters the layer.
- The minus sign is crucial: if it is missing, one gets the unphysical situation of runaway cooling or heating.

Radiative equilibrium: the working equation

Ignoring work done on the system ($DP/Dt = 0$) and atmospheric dynamics ($D/Dt = \partial/\partial t$) gives the equation typically used to solve for radiative equilibrium in a one-dimensional model,

$$\boxed{\frac{\partial T}{\partial t} = -\frac{1}{\rho c_P}\frac{\partial F_-}{\partial z}} \qquad (z \text{ increasing with altitude}).$$

- The computational iteration between the fluxes and temperatures, which depend on each other, is performed until the temperature stops changing in time (in practice until $\partial T/\partial t$ falls below a threshold, e.g. 10^{-6}; an exact zero is impossible numerically). Then radiative equilibrium is formally obtained: $\mathrm{d}F_-/\mathrm{d}z = 0$, a constant net flux through every layer.
- In frequency-resolved form the same condition reads $\int_0^\infty \kappa_\nu (J_\nu - B_\nu)\,\mathrm{d}\nu = 0$ in each layer: the energy absorbed equals the energy emitted. This is what fixes $T(\tau)$ and therefore B in the source function.
- Strictly speaking, if a calculation does not enforce radiative equilibrium, then we should not consider it a bona fide radiative transfer calculation of an atmosphere.

Local versus global energy conservation

- **Radiative equilibrium** is a condition on *every layer*: the net radiative flux through the atmosphere is the same at all heights, so no layer is heating or cooling. It determines the temperature–pressure profile.
- **Global energy conservation** is a condition on the atmosphere *as a whole*: the energy entering (absorbed starlight plus interior heat) equals the energy exiting (thermal emission to space). This is the energy balance of Lecture 2, which determines T_{eff} but not the profile.
- The two are often confused. Global conservation is a *necessary but insufficient* condition for radiative equilibrium: an atmosphere can have the right total emission yet an internally inconsistent profile in which some layers heat and others cool. A retrieval that enforces only the global condition may return such a profile.
- In the presence of **dynamics** radiative equilibrium is replaced by radiative–dynamical equilibrium: the divergence of the radiative flux is balanced by advection and by convection (where the radiative gradient exceeds the adiabatic one, Heng Chapter 4 and Seager Chapter 9). Even then the global condition holds for the planet as a whole, averaged over time.
- Heng Chapter 4 casts the distinction in precise mathematical terms using the double-grey solutions.

Clouds: the outstanding puzzle

The formation of aerosols or condensates in a gas is an outstanding puzzle of astrophysics (brown dwarfs, protoplanetary disks, exoplanets) and of Earth's climate science. Heng uses cloud, haze, aerosol, and condensate interchangeably, aware of the subtle differences: hazes are usually photochemical products; clouds follow the thermodynamic condensation curves. Forming clouds from first principles requires:

- **Nucleation:** how the precursor seed particles form, on which further growth proceeds; a true theory predicts the shape, composition, and size distribution of the particles.
- **Self-consistent gas and solid chemistry:** as the gas is depleted to form solids, its atomic and molecular abundances adjust.
- **Self-consistent opacities and radiative transfer** for gas and cloud particles, with the absorption and scattering of each computed from first principles across λ, T, and P rather than by ad hoc parameters.
- **Dynamics:** to stay aloft, the terminal velocity of a particle under gravity must be balanced by upwelling flow; sedimentation removes condensates below the photosphere. Usually parametrized by an eddy mixing coefficient K_{zz}.

Forming the particles is hard; once formed, modeling their effect on the spectrum is not controversial. "Aerosol" is the neutral term: any particle described by its size, composition, and index of refraction (or extinction efficiency).

What clouds do to a spectrum

To lowest order the aerosol adds an extinction efficiency $Q_e(x)$ with $x = 2\pi r/\lambda$ (the fitting formula and Heng Fig. 2.2 above):

- For $x \gg 1$, Q_e is roughly constant: a *grey* continuum opacity. For $x \ll 1$ it is Rayleigh-like, $\propto x^4$.
- The parameter Q_0 is a proxy for composition: refractory species (silicates) have $Q_0 \approx 10$, volatile ones (water) $Q_0 \approx 40$–80. Particle radius has a somewhat larger influence than composition.
- A cloud particle influences absorption and scattering only at $\lambda \lesssim 2\pi r$. This generalizes to a size distribution.

Consequences for spectra. When the cloud optical depth approaches unity above the gas photosphere, spectral features and lines become *muted* and less sharp for $\lambda \lesssim 2\pi r$ — the flat transmission spectra of GJ 1214b and the weakened water bands of many hot Jupiters. But the same muting can be produced by *lower* abundances of the absorbing species, so clouds introduce a **degeneracy** in the interpretation of chemical abundances. To the next order, particles are described by ω_0 and g_0, layer and wavelength dependent, in a two-stream or multi-stream calculation.

Atmospheric retrieval: letting the data talk

Traditionally, when one builds a model of an atmosphere, one starts with a governing equation — of fluid dynamics, radiation, chemistry, or some combination — and a set of assumptions, and computes. Ideally the calculation predicts a measurable quantity that may be confronted by the observations, and a judgment is made on whether the model has represented Nature — but only against the preconceived notions we have about an atmosphere.

- **Atmospheric retrieval** tries to shed those notions by providing a framework in which one *inverts* a set of observables (a spectrum) to obtain the properties of the atmosphere: its thermal structure and chemical abundances.
- Origin: inversion techniques for remote-sensing observations of Earth (Rodgers 2000), later used in the planetary sciences for the Solar System bodies. There, a combination of satellite and in-situ data provides ground truth to anchor the retrieval. "We can simply let the data do the talking."
- The output is not a single model but a *posterior distribution*: the range of atmospheres consistent with the data and their uncertainties.

Why retrieval is harder for exoplanets

- For exoplanets the faith in the data alone is less viable: we are dealing with remote sensing of *unresolved point sources*, and in-situ measurements are out of the question.
- The data are often sparse and fragmented, in the sense that one only sees incomplete facets of the atmosphere — the dayside but not the nightside, insufficient wavelength coverage, a handful of photometric bands — that do not allow unique interpretations.
- Letting the interpretation be guided entirely by the data, for example without enforcing radiative or chemical equilibrium, becomes foolhardy when the number of model parameters exceeds the number of data points. Our interpretation needs to be anchored by the laws of physics and chemistry to some degree.
- How much physics to impose is an ongoing debate: "free" retrievals with parametrized profiles and abundances versus "self-consistent" models with equilibrium chemistry and radiative equilibrium. The former are flexible but can wander into unphysical corners; the latter are constrained but may be wrong if the assumed physics is incomplete.
- JWST-quality spectra, with hundreds of spectral points, have shifted the balance: the data now have enough information to test the assumptions of the self-consistent models rather than merely to fit a few parameters.

Basic ingredients of a retrieval model

Augmenting the radiative-transfer checklist:

1. A way of describing the **temperature–pressure profile** with a set of parameters: either an ad hoc fitting function, or a simplified analytic solution of the radiative transfer equation (Heng Chapter 4, the double-grey solutions).
2. The **identities of the atoms and molecules** included (e.g. water, methane, carbon monoxide). Their relative abundances may be free parameters of the retrieval; alternatively one may enforce *chemical equilibrium*, in which case the molecular abundances are determined by the elemental abundances and only the latter are free (Heng Chapter 7).
3. For a fixed combination of parameter values, a **forward model** converts the opacities into a synthetic spectrum. This conversion is performed once for each parameter set; there is typically no attempt to solve for radiative equilibrium, as the iteration described earlier is not performed.

Scanning the parameter space

- Millions, if not billions, of combinations of these parameters are considered. Each synthetic spectrum is compared against the observed one and a goodness of fit (a likelihood, typically $\propto e^{-\chi^2/2}$) is computed for each model.
- In this manner a range of best-fit models is found, given the uncertainties on the data, and the *posterior distributions* of the molecular abundances and of the temperature profile may be computed. Correlations between parameters (the degeneracies of the next slides) show up as elongated or banana-shaped posteriors.
- In practice a minimization or sampling algorithm is required to scan this large, multidimensional parameter space efficiently and thoroughly:
  - **Optimal estimation** exploits the property that if the prior of an input quantity is Gaussian, the expressions used for minimization simplify (Rodgers 2000). A uniform prior is mimicked by a very wide Gaussian.
  - **Monte Carlo methods** (Markov-chain Monte Carlo, nested sampling) allow general priors and have been implemented in various flavors; they are now standard in exoplanet retrievals.
- Because the forward model is evaluated millions of times, its speed is critical: this is why retrievals use fast layer-by-layer pure-absorption solutions with precomputed opacities rather than full radiative-equilibrium iterations.

Retrieval pitfalls (1)

A retrieval model is only as good as the assumptions and techniques used. Some of the possible pitfalls:

- **Degeneracies.** The spectral lines of different molecules inevitably overlap across wavelength, so multiple combinations of relative abundances may lead to the same net absorption at a given wavelength. Aerosols introduce another degeneracy: they diminish the strength of spectral lines, an effect that may be compensated by increasing the abundances of the molecules responsible. Any good retrieval quantifies these degeneracies formally and systematically (correlated posteriors).
- **Being misled by your priors.** The posteriors depend on the priors assumed. Example: the carbon-to-oxygen ratio estimated as

$$\mathrm{C/O} \approx \frac{\tilde{n}_{\mathrm{CO}} + \tilde{n}_{\mathrm{CH_4}} + \tilde{n}_{\mathrm{CO_2}}}{\tilde{n}_{\mathrm{CO}} + \tilde{n}_{\mathrm{H_2O}} + 2\tilde{n}_{\mathrm{CO_2}}},$$

  with $\tilde{n}$ the mixing ratios (abundances normalized by H_2). If CO dominates, C/O ≈ 1; if CO_2 dominates, C/O ≈ 0.5; if CH_4 or H_2O dominate, C/O $\gg 1$ or $\ll 1$. So even uniform priors on the molecules produce two artificial peaks in the prior distribution of C/O.

Retrieval pitfalls (2)

- **Local versus global energy conservation.** In the ensemble of models it is natural to discard those whose emergent thermal emission exceeds the energy input from the star (for negligible internal heat). That is simply global energy conservation: you cannot produce more than what you put in. What is less obvious is *local* energy conservation, radiative equilibrium: the net heating or cooling of each layer must be the same across layers, in the absence of circulation, and most retrievals do not enforce it.
- **Not everything is chemically possible.** In chemical equilibrium not every combination of relative abundances is possible: each layer has its own T and P, and the mixing ratios take specific values given the elemental abundances. One is not free to specify them as independent parameters. When dynamical mixing is faster than chemistry the mixing ratios depart from local equilibrium — but again not arbitrarily: with the *quenching* approximation (Heng Chapter 6) they originate from a deeper layer that *is* in equilibrium. Example: CO_2 in hot ($\gtrsim 1000\,\mathrm{K}$) hydrogen-dominated atmospheres is never more abundant than CO or H_2O unless the metallicity is implausibly high.

Retrieval pitfalls (3)

- **Opacities: the devil is in the details.** Across the UV to infrared one must compute the shapes of millions to billions of lines, for one molecule at one (T,P). Line-by-line calculations sample all of them without approximation, a formidable task at $\sim 1000\,\mathrm{K}$; binning techniques such as the k-distribution method speed this up at the cost of approximations (Heng Chapter 5). Hidden in the details is the pressure broadening of the far line wings, an unsolved physics problem. And photometric observations integrate over broad bandpasses, so many combinations of molecules can produce the same opacity across the band. When published papers omit these details, exercise skepticism — one of the most precious instruments in a scientist's toolkit.
- **Numerical convergence.** Running a retrieval requires discretizing space (the T–P profile), wavelength (the opacity function), and the set of species; to finish in reasonable time one often uses coarse resolution along one axis. Determining whether the resolution is sufficient requires resolution tests — thankless but necessary. As a rule of thumb (attributed by Heng to Scott Tremaine), always adopt the second-highest resolution if the highest one attains convergence.

Take-home messages

Flux ratios. $F_{\oplus,p}/F_{\oplus,*} = F_{\mathrm{S},p}R_p^2/F_{\mathrm{S},*}R_*^2$. Thermal: $(R_p/a)^2(f/4)(1-A_{\mathrm{B}})$ bolometrically, $(T_{\mathrm{b},p}/T_{\mathrm{b},*})(R_p/R_*)^2$ in the Rayleigh–Jeans regime. Reflected at full phase: $A_g R_p^2/a^2$ (geometric, not Bond, albedo; $g = 4A_g/A_{\mathrm{S}}$). Transmission: $2R_pA_H/R_*^2 \sim 10^{-3}$ for hot Jupiters. Phase curves: $A_g(R_p/a)^2\Phi_\alpha$ with the Lambert $\Phi_\alpha = [\sin\alpha + (\pi-\alpha)\cos\alpha]/\pi$; thermal phase curves follow the same law for instantaneous reradiation.

Optical depth. $\tau = \int n\sigma\,\mathrm{d}x = \int \kappa\rho\,\mathrm{d}x = \int \alpha_e\,\mathrm{d}x$; photons travel to $\tau \sim 1$; this defines photospheres, transit radii ($\tau_{\mathrm{chord}} = n\sigma\sqrt{2\pi HR}$), and regimes (streaming versus diffusion, $N \sim \tau^2$). Mean free path $1/\kappa$.

Opacity and emission. $\kappa = \alpha + \sigma_s$ (isotropic); $\varepsilon = \alpha B + \sigma_s \frac{1}{4\pi}\oint PI\,\mathrm{d}\Omega'$; particles: $\sigma = Q_e\pi r^2$, ω_0, g_0; Planck and Rosseland means for thin and thick regimes. LTE: matter properties depend only on local T and P (collisions dominate); Kirchhoff $\varepsilon = \alpha B$; source function $S = (1-\omega_0)B + \omega_0 J$.

Transfer equation. $\mu\,\mathrm{d}I/\mathrm{d}\tau = I - S$ in plane-parallel geometry; formal solution $I(0,\mu) = \frac{1}{\mu}\int_0^\infty S e^{-\tau/\mu}\mathrm{d}\tau$; hard because S depends on I (scattering, radiative equilibrium) and the boundary conditions are split between top and bottom. Beer's law for pure extinction; isothermal layers exchange flux with $\mathscr{T} = 2E_3(\Delta\tau)$. Radiative equilibrium is local, not just global, energy conservation. Clouds add grey extinction for $\lambda \lesssim 2\pi r$ and degeneracies; retrievals invert spectra but must respect physics.

References

- Seager, S. 2010, *Exoplanet Atmospheres: Physical Processes* (Princeton University Press), Sections 3.5–3.6 and Chapter 5.
- Heng, K. 2017, *Exoplanetary Atmospheres: Theoretical Concepts and Foundations* (Princeton University Press), Chapter 2.
- Mihalas, D. 1978, *Stellar Atmospheres*, 2nd ed. (Freeman) — the most thorough treatment of LTE.
- Rybicki, G. B. and Lightman, A. P. 1986, *Radiative Processes in Astrophysics* (Wiley), Chapter 1.
- Chandrasekhar, S. 1960, *Radiative Transfer* (Dover).
- Goody, R. M. and Yung, Y. L. 1989, *Atmospheric Radiation: Theoretical Basis*, 2nd ed. (Oxford).
- Pierrehumbert, R. T. 2010, *Principles of Planetary Climate* (Cambridge).
- Sobolev, V. V. 1975, *Light Scattering in Planetary Atmospheres* (Pergamon).
- Cox, A. N. 2000, *Allen's Astrophysical Quantities*, 4th ed. (Springer) — phase-curve formulae.
- Toon, O. B., et al. 1989, *JGR* **94**, 16287 (two-stream with linear source function).
- Rodgers, C. D. 2000, *Inverse Methods for Atmospheric Sounding* (World Scientific) — optimal estimation.
- Madhusudhan, N. and Seager, S. 2009, *ApJ* **707**, 24; Line, M. R., et al. 2013, *ApJ* **775**, 137; Benneke, B. and Seager, S. 2012, *ApJ* **753**, 100 — exoplanet retrieval.
- Lecavelier des Etangs, A., et al. 2008, *A&A* **481**, L83; Fortney, J. J. 2005, *MNRAS* **364**, 649 — chord optical depth and transit radius.