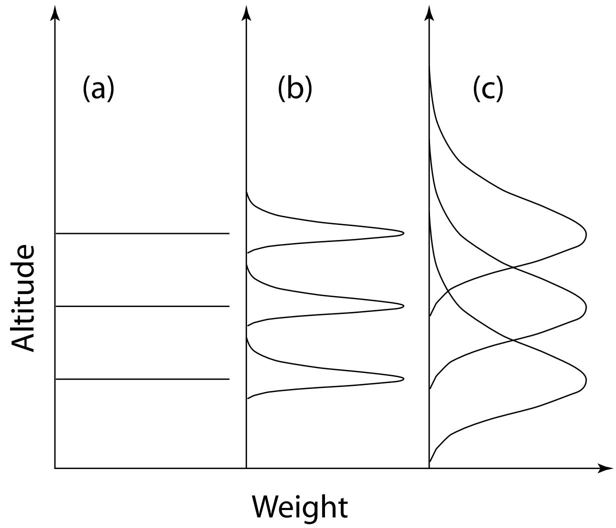

Fig. 8.4: Idealized satellite sensor weighting functions for atmospheric temperature profile retrievals. (a) Channel weighting functions resemble δ-functions; i.e., all emission observed by each channel originates at a single altitude. (b) Observed emission represents layer averages, but channel weighting functions do not overlap. (c) Realistic case, in which weighting functions not only represent layer averages but also overlap.

on a rigorous discussion of remote sensing theory, which is best left for a separate course and / or textbook. It is enough for now that you recognize the close connection between the radiative transfer principles discussed earlier in this chapter and an application of immense practical importance to modern meteorology.

Figure 8.4 depicts the physical basis for profile retrieval at three levels of idealization, starting with the simplest — and least realistic — on the left: If weighting functions happened to be perfectly sharp — that is, if all emission observed at each wavelength λ_i originated at a single altitude (Fig. 8.4a), then "inverting" the observations would be simple: in this case, the observed brightness temperatures $T_{B,i}$ would exactly correspond to the physical temperatures at the corresponding altitudes h_i. Your job is then essentially finished without even lifting a calculator. Of course, you wouldn't know how the temperature was varying *between* those levels, but you could either interpolate between the known levels and hope for the best or, if your budget was big enough, you could add an arbitrary number of new channels to your sensor to fill in the vertical gaps.

Slightly more realistically, panel (b) depicts the weighting functions as having finite width, so that the observed brightness temperatures correspond to an *average* of $B_\lambda[T(z)]$ over a substantial depth of the atmosphere rather than a unique temperature T_i at altitude h_i. There is now ambiguity in the retrieval, because there is no single atmospheric level that is responsible for all of the emission measured by any given channel. At best, you can estimate the *average layer temperature* associated with each channel. Nevertheless, the profile retrieval problem itself remains straightforward, because each channel contains completely independent information: there is no overlap between the weighting functions.

Unfortunately, real weighting functions are constrained to obey the laws of physics, as embodied in (8.30). This means that unless you have an unusually sharp change with altitude in the atmospheric absorption coefficient β_a, your weighting functions will be quite broad. In the worst case, the mass extinction coefficient k_a of your chosen constituent will be nearly constant with height, so that the weighting functions will be essentially those predicted for an exponential absorption profile as discussed in Section 7.4.3. The situation is somewhat better for sensor channels positioned on the edge of an absorption line (or between two lines), because *pressure broadening* (see Chapter 9) then increases k_a toward the surface, which sharpens the weighting function. Nevertheless, the improvement is not spectacular.

Therefore, given any reasonable number of channels, there is always considerable overlap between adjacent weighting functions, as depicted schematically in Fig. 8.4c. In fact, Fig. 8.5 shows actual weighting functions for the Advanced Microwave Sounding Unit (AMSU), which has 11 channels on the edge of the strong O_2 absorption band near 60 GHz (c.f. Fig. 7.7). Although each satellite sounding device has its own set of channels and therefore its own unique set of weighting functions, those for the AMSU are fairly typical for most current-generation temperature sounders in the infrared and microwave bands.

To summarize: It is clear on the one hand that there is information about vertical temperature structure in the radiant intensities

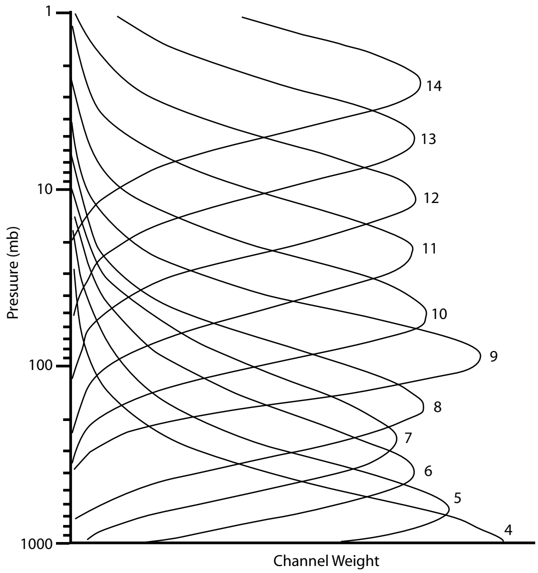

Fig. 8.5: Weighting functions for channels 4–14 of the Advanced Microwave Sounding Unit (AMSU).

observed by a sounding instrument like the AMSU. On the other hand, one shouldn't underestimate the technical challenge of retrieving temperature profiles of *consistently useful quality* from satellite observations. In outline form, here are the main issues:

- In general, it takes far more variables to accurately describe an arbitrary temperature profile $T(z)$ than there are channels on a typical satellite sounding unit. This means that you have fewer measurements than unknowns, and the retrieval problem is *underdetermined* (or *ill-posed*). The problem is therefore not just that of finding any temperature profile that is consistent with the measurements; the real problem is of choosing the *most plausible* one out of an infinity of physically admissible candidates.

- Because of the high degree of vertical overlap between adjacent weighting functions, the temperature information contained in each channel is not completely independent from that provided by the other channels. That is to say, if you have