Chapter 1

Introduction and overview

1.1 VELOCITY ANALYSIS FOR REFLECTION SEISMOLOGY

Seismic reflection experiments comprise initiation of waves in the Earth using artificial energy sources at the surface, and recording of the resulting wavefields at other locations on the surface. These waves propagate through rocks that are elastic, heterogeneous in both density and wave velocity on many scales, anisotropic, dispersive, and attenuating. Nonetheless, much of what is seen in seismic reflection data often can be explained as primary reflections in a medium whose physical properties are largely controlled by variations in a single parameter, the compressional wave (P-wave) velocity. Other effects certainly contribute to the data, but the P-wave primaries are usually sufficiently dominant that conventional processing uses only this information, leaving the extraction of information about other parameters to specialized, subsequent processing; good knowledge of the P-wave velocities is usually required as a prerequisite for more sophisticated analysis techniques.

Figure 1.1 shows a small portion of an image obtained from processing a marine seismic reflection survey. Figure 1.2 is a sample of the field data before most processing, showing the suite of records from a single firing of a watergun energy source. These data are recorded by a linear cable of hydrophones that sense fluctuations in pressure, so they represent a band-limited record of a wavefield in time. After processing, the data still resemble a wavefield of some sort, but are now readily interpretable to the human eye as a picture of complexly folded subsurface geological strata. These processed data represent an approximate image of the subsurface structure, with high amplitude events occurring where the juxtaposition of rocks with different material properties causes discontinuities in the acoustic velocities. Because the original data are bandlimited the processed image will likewise appear bandlimited. The low wavenumbers are missing in Figure 1.1; one cannot see the trends in the velocity, which in these data, as in most, increases steadily with depth.

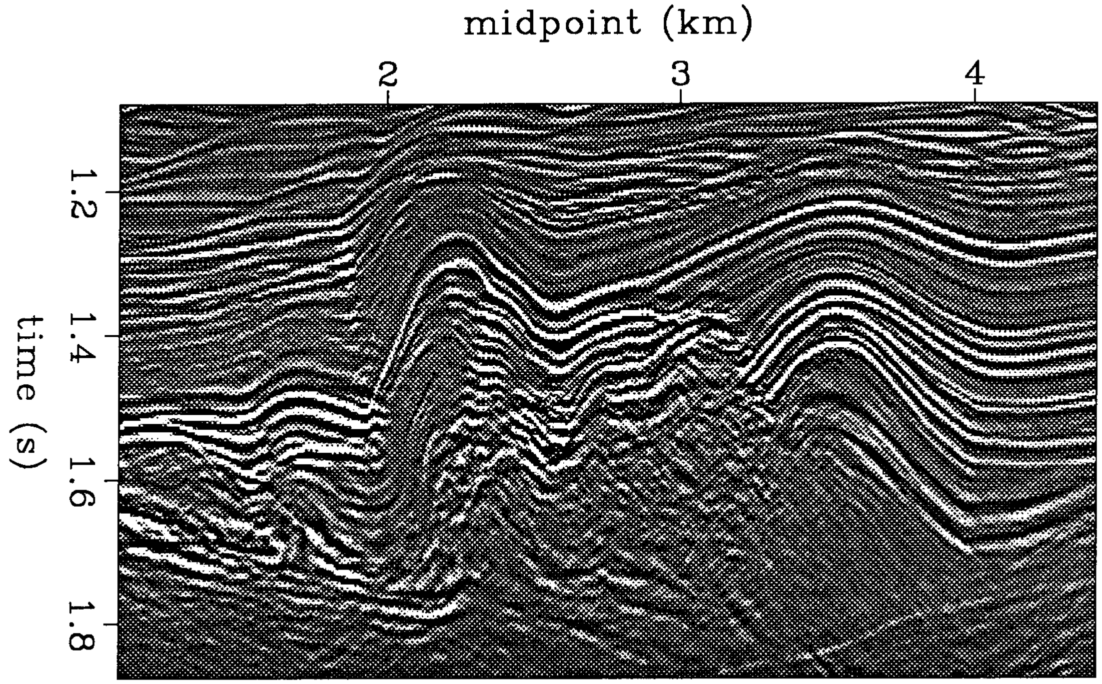

FIG. 1.1. Portion of a migrated marine survey. These data are from offshore southern California, and are analyzed more fully in chapter 2.

The problem of obtaining an image of the Earth's subsurface from seismic reflection data is conventionally divided into a velocity analysis stage, followed by an imaging, or migration, stage. Figure 1.1 shows data after migration; an estimate of the background, or low wavenumber, part of the velocity field has to be made before conventional migration imaging, and then provided as additional input to the computer program that implements the migration algorithm. The low wavenumber information using for migrating the data in Figure 1.1 is shown in a separate form as a background velocity field in Figure 1.3. This information is also derived from the recorded data, but by very different methods than used to produce Figure 1.1. This conventional dichotomy reflects the understanding that the problem of finding the P-wave velocities must be treated differently for different spatial wavenumber domains. The background velocity field is derived principally from examining what happens to the waves as they are *transmitted* through the rocks; the migrated image is concerned with where the waves are *reflected* from interfaces. The coherent reflections arise from velocity changes that are high wavenumber, that is, that are rapid relative to the length scale set by the bandwidth of the acoustic waves

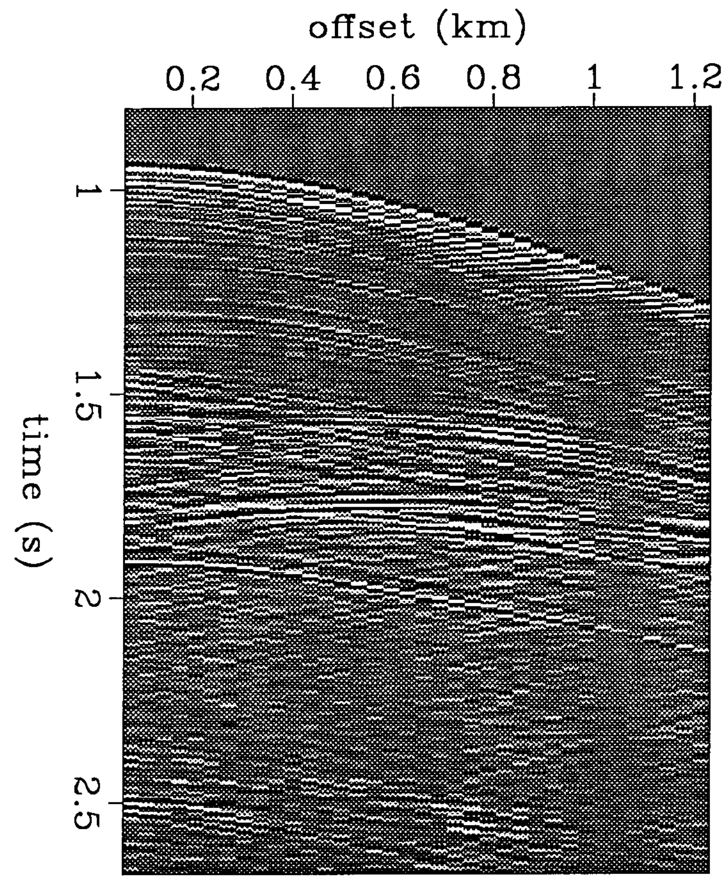

FIG. 1.2. Shot profile from a marine survey. This is an example of the field data that were used in producing Figure 1.1.

traveling across them. In practice, high wavenumber images are usually parametrized not by the velocity itself, but by the reflectivity, which is more closely related to the first derivative of the velocity. Migration is used to convert the observed wavefields from functions of recording time and position, into images of the reflectivity as a function of the subsurface location and depth. However, migration requires as input the specification of a velocity field. This is not as problematic as it might appear, because the velocity information needed is primarily the low wavenumber components of the velocity field, not the high ones. The observed positions of the reflectors in the data, that is, the traveltimes to events, depend on integrals, or averages, of the velocity field seen by the transmitted wave as it passes through the overlying rocks. Such an integral is most sensitive to the low wavenumber components of the velocity, unlike the reflectivity, which depends on the derivatives of the velocity, and is thus most sensitive to high wavenumber components.

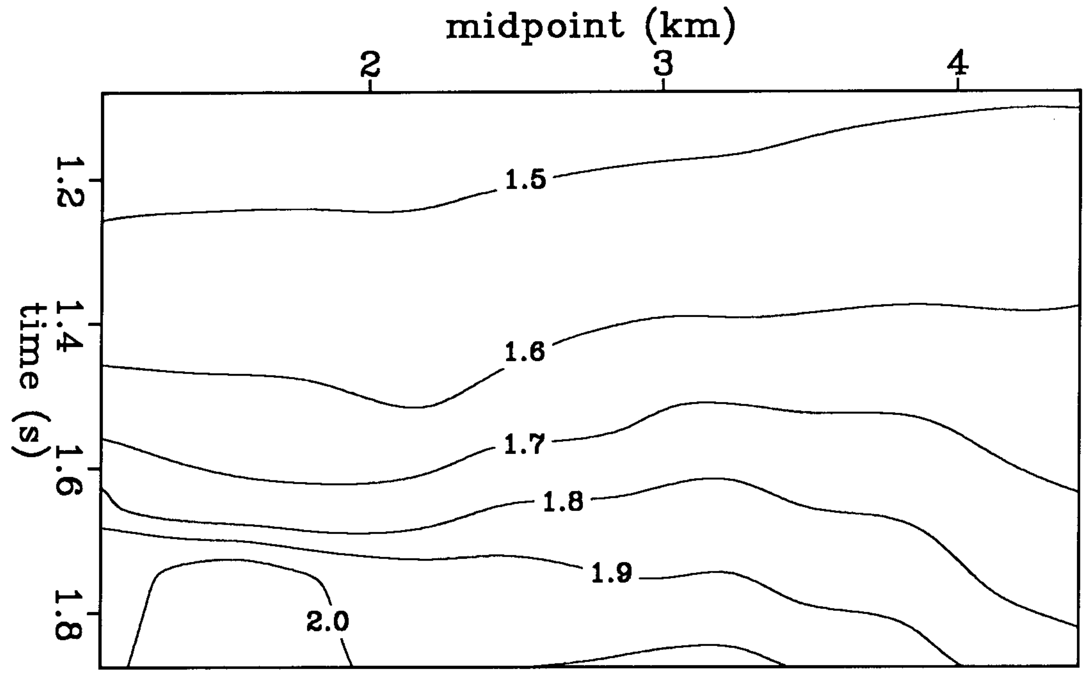

FIG. 1.3. Velocity field used for producing Figure 1.1. The background migration velocity field is shown contoured in intervals of 0.1 km/s.

The decoupling of the inversion of seismic wavefields for P-wave velocity into high and low wavenumber domains is necessarily imperfect. There is no simple cutoff wavenumber at which one method replaces another; the wavenumber content of the data is determined both by the experiment and by the geologic structure present. Moreover, one needs to know the reflector locations to invert traveltime information for the background velocity field, but one needs the velocity field to migrate the data and find the reflector locations. To escape this apparent paradox, one needs to introduce enough simplifying assumptions or a priori information to get an estimate of one or the other, and then iteratively bootstrap to a better answer. The most common way that additional information is introduced is by picking, either of traveltimes, or of stacking velocities.

1.2 STACKING VELOCITY ANALYSIS

The low wavenumber velocity information is contained in the traveltimes, so methods that start with the explicit picking of traveltimes to events in the data are perhaps conceptually the most straightforward. In practice, however, methods that use picking of stacking velocities are in more widespread use, so I will defer until the next section discussion of methods that use traveltimes directly.

For velocity analysis, the data are usually sorted into gathers of all traces that share a common midpoint (CMP). For such a CMP gather, in a constant velocity medium, the reflections from a flat, horizontal layer satisfy a simple hyperbolic moveout equation,

$$t^2 = t_0^2 + \frac{x^2}{v^2} \tag{1.1}$$

where t is the traveltime, x is the lateral distance between source and receiver, v is the velocity, and t_0 is the traveltime for the vertically incident ray corresponding to coincident source and receiver. The next simplest medium to consider is one with a series of flat, stratified layers. The moveout will no longer be exactly hyperbolic, but if the velocity variations are moderate, the moveout can be approximated well by equation 1.1 provided the velocity v is interpreted as a root-mean-square (rms) average of the velocities in the overlying layers. Estimation of rms velocities in this way, and inversion of them to find the interval velocities in each layer, is commonly called "the Dix approximation," after the pioneering work of Dix (1955).

To estimate velocities, then, one attempts to fit moveout hyperbolas to events in CMP gathers. This can be done in a variety of ways, the most common of which is to apply an NMO correction according to equation 1.1, using a range of trial velocities, and then to evaluate a statistic such as the mean or the semblance over offset for each event, seeking to maximize the resulting value. Evaluating the mean corresponds to forming repeated stacks of the data for various velocities; for flat beds, NMO stacking provides an adequate imaging operator. Such an analysis yields apparent stacking velocity as a function of zero-offset traveltime; the interval velocity can be found as a function of depth by inversion using the rms assumption.

Even in a constant velocity medium, however, this method of velocity analysis runs into problems if the reflectors are not horizontal. The moveout will remain hyperbolic, but with an apparent velocity increased by the secant of the dip angle. The rms assumption will break down, and when events in the seismic data cross, the velocity function will become multi-valued. An additional processing step, commonly called dip-moveout, or DMO, (e.g., Deregowski, 1986) can be introduced to ameliorate these problems. The drawback, however, is that DMO can be computationally expensive to apply, and it does not commute with NMO, so it must be applied once for each velocity stack created.

DMO does not solve all the problems of velocity analysis, even where the velocity depends only on the depth. A dipping event moves updip during migration; before migration the velocity information from such an event is mispositioned vertically and laterally. For this reason, it is desirable to migrate data *before* velocity analysis, except, of course, that migration requires prior knowledge of the velocities, the apparent paradox pointed out in the previous section.

The velocity gradient is generally dominantly vertical, especially in the sedimentary basins typically explored for hydrocarbons using reflection seismology. The gradient, however, can also contain a substantial lateral component, particularly in regions of complex geological structure. In such areas, velocity analysis based on NMO encounters further problems, because moveout can deviate significantly from hyperbolic patterns. Often it is still possible to find an approximate, best-fitting hyperbola, but the assumption that the apparent velocity is an rms average is no longer adequate. Instead, one needs to unscramble the relations between the observed laterally varying stacking "velocities" and the underlying physically significant interval velocities. This problem has been successfully addressed for the restricted case of flat beds (Loinger, 1983; Toldi, 1985).

Stacking velocity analysis requires picking the apparent velocities of major events in the data. This picking can be difficult, because the desired peaks can be ill-defined and buried in a sea of noisy subsidiary maxima. Moreover, small changes in stacking velocity picks can imply large nonphysical changes in the implied interval velocities; the inversion is *unstable*. Toldi (1985) has shown how this problem of instability can be lessened by formulating the inversion as a constrained and damped optimization problem, avoiding explicit picking of peaks. However, even if explicit picking is used, the enhancement of signal-to-noise ratio effected by stacking makes the accurate picking of stacking velocity peaks easier than the direct picking of traveltimes in prestack data required by traveltime tomography, which I consider next.

1.3 REFLECTION TOMOGRAPHY

Inversion based on picked traveltimes is generally termed *tomography*. The word is from the Greek $\tau o\mu\eta$, meaning a cut or slice, and $\gamma\rho\alpha\phi\omega$, to write; it refers to any reconstruction of properties of an object from values measured along line or surface integrals cutting through the object. Tomography is perhaps most familiar in medical usage, where the development of computerized tomographic imaging of density contrasts in the human body from axial X-ray scans won Godfrey Hounsfield and Allan Cormack the 1979 Nobel prize in physiology or medicine. In reflection seismology it is usually the velocity distribution one wishes to reconstruct; the traveltimes used as data are line integrals of the slowness (reciprocal of the velocity) along the ray paths.

Many papers have been written describing seismic tomographic methods; I shall mention only a few here. (For a wide-ranging recent collection of papers, see Nolet, 1987). One of the most representative and thorough investigations is that described by Bishop et al. (1985). In their approach, the goal is to construct a model of reflector positions and interval slownesses so that traveltimes computed by ray tracing for selected events match as closely as possible traveltimes that are picked from the seismic data. Similar methods are discussed by Gray and Golden (1983) and by Chiu et al. (1986). Another closely related tomographic method presented by Sword (1987) incorporates picked estimates of ray parameters as well as traveltimes. All these approaches result in nonlinear, iterative, least-squares algorithms.

Tomographic velocity analysis requires either prior knowledge of reflector position, or simultaneous estimation of these positions. Stork and Clayton (1986) proposed combining a tomographic solution for velocities with iterative depth migration for locating reflectors; in their scheme, the reflector locations can be picked anew with each update of the velocity model. This hybrid method fits well with the idea that resolution of high and low wavenumber components of the model require different techniques, and using a separate migration step avoids many instability problems that can arise from inverting traveltime picks for reflector locations as well as for velocities.

Perhaps the greatest drawback of traveltime tomography is the necessity for large amounts of picking. Events in seismic data are generally complicated and variable wave patterns; reducing this information to isolated time picks can involve large amounts of human judgement and can be very time consuming. The picking process is subject to systematic errors caused by ambiguities in defining events or by incorrect assumptions about wavelet phase, as well as to random errors. Automatic picking programs can be faster than people, but human judgement is usually better at avoiding egregious mispicks. Picking in any form is an irreversible process that inherently oversimplifies the potentially exploitable information in the data.

Because of the difficulties in picking, seismic tomography rarely uses information from more than a few distinct horizons. The limited aperture (maximum offset) of the reflection experiment often severely limits the vertical resolution of velocity anomalies; only an average slowness between reflectors can be determined (see, e.g., Bishop et al., 1985, or for a more thorough discussion of resolution problems, see Stork, 1988). Improving resolution requires using more intermediate reflectors, and consequently more picking.

Many tomographic and quasi-tomographic methods use models of constant velocity blocks bounded by reflecting horizons (e.g., Thorson et al., 1987). Such models are easy to construct, and easy to trace rays through. However, they imply that the high wavenumbers (reflectors) are tightly coupled to the low wavenumber velocity trends (by insisting that velocities change only as a few abrupt steps), an assumption that is too limiting to represent many real geological conditions. More sophisticated models allow for laterally varying velocities within blocks (e.g., van der Made et al., 1987) or for definition of velocity fields on grids distinct from those used for defining reflector locations (e.g., Bishop et al., 1985, Sword, 1987, or Stork, 1988). I use here only the last type of velocity model, in which the background velocities used for computing transmission effects are decoupled from the reflectivity image.

1.4 VELOCITY ANALYSIS USING PRESTACK MIGRATION

Conventional stacking velocity analysis is predicated on the idea that the best velocity field is the one for which NMO correction makes the traces at different offsets most similar. Its biggest failings are caused by the inability of NMO alone to image all the data adequately. It is inviting to try to extend the principle to a method that would use a more sophisticated imaging operator such as prestack migration. In complex structural areas, depth migration is usually indicated. Such a velocity analysis scheme, using shot-profile depth migration, has been suggested by Al Yahya (1987). The largest hurdle in implementing such an algorithm arises from the type of velocity field used, and the size of the implied model parameter space that must be searched. NMO uses a single parameter, the moveout velocity, to describe the imaging operator at each point. Depth migration, however, uses at every point the interval velocities for all the overlying points, increasing the number of parameters by orders of magnitude. That is, to go from one NMO stack to a better one, only a single number, the stacking velocity, has to be changed. This allows for the evaluation of a stacking velocity function by a simple scan over a range of stacking velocities, followed by a later inversion of the stacking velocities for interval velocities. In contrast, going from one depth migration to a better one requires changing the interval velocities at many points, which cannot be done practically by a simple exhaustive search.

Is there any way to use prestack migration for velocity analysis and still retain the simplicity and ease of NMO methods? Al Yahya (1987) presented one possible approach, using a single parameter search to reduce residual curvature in migrated shot profiles. Here I take a different tack, one that directly extends conventional analysis by replacing NMO as an imaging operator not by *depth* migration, but by prestack *time* migration. The NMO stacking operator images data by summing over moveout hyperbolic curves defined by equation (1.1). Imaging using prestack time migration conceptually involves a similar data summation, this time over the diffraction traveltime surface given by

$$t = \left(\frac{t_0^2}{4} + \frac{(x_s - x_0)^2}{v^2} \right)^{1/2} + \left(\frac{t_0^2}{4} + \frac{(x_g - x_0)^2}{v^2} \right)^{1/2} \tag{1.2}$$

where x_s is the shot location, x_g is the geophone location, and x_0 is the location of the image point. Like the NMO stacking operator, prestack time migration depends on a single parameter v, the apparent velocity, that controls the curvature of the operator. Also like NMO stacking, it is a correct imaging operator for constant velocity media, but now it treats all dips, not just horizontal beds.

Hale (1984) showed that prestack time migration is equivalent to applying the sequence of NMO, DMO, stack, and zero-offset migration. In chapter 2, I show how DMO and migration can be implemented on a sequence of constant velocity stacks, thus converting conventional velocity analysis using scans over stacking velocities into a more sophisticated velocity analysis using scans over DMO-corrected stacking velocities or over prestack time-migration velocities. These last two types of velocity analyses extract the same information from the data, and differ principally in the positions (migrated or unmigrated) at which the velocities are evaluated. For the discussion here, I often use the phrase "prestack time migration" generically to refer to both velocity analysis methods. I discuss the practical differences further in chapter 4; my preliminary results there suggest that full prestack time-migration is more useful for interactive analyses based on focusing, whereas DMO-corrected stacks work better for automatic methods based on maximizing total energy.

Applying any of these imaging operators (NMO stacking, DMO-corrected stacking, or prestack time migration), over a range of velocities, replaces the offset dependence of the data with a velocity parameter. This transformation is linear, and if enough velocities are used, the transformation is nearly invertible. (See Thorson, 1984, or Thorson and Claerbout, 1985, for an extensive discussion of the invertibility of velocity stacking.) Thus, little information is lost in the transformation to such a velocity space. The advantage gained is that, provided the moveout in the original data is roughly hyperbolic, the useful parts of the data will be clustered or compacted in the new velocity space. The velocity transformation emphasizes the velocity information so that peaks can be picked, and a simple curve through these peaks will describe most of the gross properties of the velocities.

Prestack time migration, like NMO stacking, works best for laterally invariant velocities. Imaging and velocity analysis still work for a fair degree of lateral velocity variation, but the "velocities" found will not be simply related to the underlying interval velocities. As mentioned above, a method for inverting stacking velocities is presented by Toldi (1985). In chapters 3 and 4, I discuss how to extend this approach from stacking velocities to prestack time migration velocities, removing the restriction to horizontal bedding and allowing inversion for interval velocities in complex structure. This inversion treats the migration velocities as functions of the traveltimes for different offset experiments. Standard tomography inverts these traveltimes directly. Because I invert the migration velocities instead, the inversion scheme resembles tomography with an extra filtering step included to convert traveltimes into migration velocities. This extra step obviates the need to pick events in the original data, replacing it with the picking of velocities, which I believe can be more robustly incorporated into an automatic optimization algorithm.

1.5 ASSUMPTIONS AND LIMITATIONS

Prestack time migration requires that moveout and diffraction time surfaces be approximately hyperbolic. When lateral velocity variation becomes too severe, this assumption will break down. The basic requirement for velocity analysis using time migration to work is that a unique velocity value can be found for which equation 1.2 best fits diffraction time surfaces in the data, and hence produces the strongest amplitude and most coherent images in migration. If there is any velocity variation at all, equation 1.2 will not fit perfectly, and the migration velocities will not exactly equal the medium velocities. However, a large amount of misfit can be tolerated and one can still invert for interval velocities as long as the migration velocity function remains single valued.

Fortunately, for many data, prestack time migration yields a remarkably coherent image even with substantial velocity variation. Because the diffraction surface used is always of the form given in equation 1.2, no allowance is made for the finer details of ray bending, so subsequent depth migration using an accurate interval velocity model should be able to improve the image. The principle effect, however, will often not be a substantially better focusing of events, but rather mostly just a repositioning of events. This effective decoupling of depth migration into a time migration followed by what is known as an "image ray" correction has been exploited previously for post-stack migration by Hubral (1977) and by Larner et al. (1981).

The type of data for which prestack time migration would not be expected to work well for velocity analysis is thus that for which lateral velocity variation is so strong that only depth migration can give any coherent image. Examples might include rugged sea floor topography, or overthrust faults with extreme velocity contrasts such as sometimes found in the Rocky Mountains. Good imaging of reflectors below such radical velocity contrasts can be extremely difficult. Often, however, such velocity problems are restricted to particular depths or layers. In such cases, it might be possible to solve for velocities in the offending layer, continue the rest of the data downward past the problem zone, and reduce the problems of velocity analysis and imaging in lower layers to a more manageable problem. (See Yilmaz and Chambers, 1984, for one such layer removal approach.)

Prestack migration velocity analysis will not be able to resolve every feature of the velocity field. I examine resolution in more detail in chapter 4, but some general features can be predicted in advance. Like any tomographic method using traveltime information to derive velocities, this approach will not be able to resolve all ambiguities between velocities and depths to reflectors. It will be better able to find long wavelength (low wavenumber) components of the velocity field than short wavelength (high wavenumber) ones. Fortunately, a smooth background velocity model is exactly what is usually wanted for subsequent migration. Smoothing of velocity fields normally results in little degradation of migrated images if done carefully, and inclusion of high wavenumber components in a migration velocity field can be dangerous unless they are known precisely. Finally, like many nonlinear inversion schemes, there exists no guarantee of convergence if the starting guess is poor.