Chapter 1

Kirchhoff migration with band-limited traveltimes

In this chapter I compare the images created by Kirchhoff migration using Green's functions computed by three different methods.

- First arrival traveltimes from a finite difference solution to the eikonal equation.
- Maximum amplitude arrival traveltimes, amplitudes and phases from paraxial raytracing.
- Maximum energy traveltimes amplitudes and phases computed by my band-limited Green's function method.

1.1 Kirchhoff migration in the time domain

The modeling equation for a scalar scattering function was previously stated in the frequency domain.

$$P(g,s,\omega) = \int\int G(s,x,z,\omega)G(x,g,z,\omega)S(x,z)\ dx\ dz.$$

If we wish to use this method in the time domain it must be inverse Fourier transformed over frequency:

$$P(g,s,t) = \int\int\int G(s,x,z,\omega)G(x,g,z,\omega)S(x,z)\ dx\ dz\ e^{-i\omega t}\ d\omega\,. \qquad (1.1)$$

The approximate, single-event form for the combination of the two Green's functions can be written as:

$$G(s,x,z,\omega)G(x,g,z,\omega) = A(s,x,z)A(x,g,z)\, e^{i\phi(s,x,z)}\, e^{i\phi(x,g,z)}\, e^{i\omega\tau(s,x,z)}\, e^{i\omega\tau(x,g,z)}$$

$$= A(s,x,z)A(x,g,z)\, e^{i[\phi(s,x,z)+\phi(x,g,z)]}\, e^{i\omega[\tau(s,x,z)+\tau(x,g,z)]}.$$

If this form is inserted in equation (1.1) we obtain,

$$P(g,s,t) =$$

$$\int\int A(s,x,z)A(x,g,z)e^{i[[\phi(s,x,z)+\phi(x,g,z)]}S(x,z)\int_{-\infty}^{+\infty} e^{i\omega[\tau(s,x,z)+\tau(x,g,z)]}e^{-i\omega t}\ d\omega\ dx\ dz$$

$$= \int\int A(s,x,z)A(x,g,z)S(x,z)\, e^{i[\phi(s,x,z)+\phi(x,g,z)]}\delta(t-[\tau(s,x,z)+\tau(x,g,z)])\ dx\ dz.$$

This equation describes an integral form of prestack modeling in the time domain. For each subsurface point there is a single arrival on the surface seismogram. It arrives at time $\tau(s,x,z)+\tau(x,g,z)$ with a phase of $\phi(s,x,z)+\phi(x,g,z)$ and an amplitude of $A(s,x,z)A(x,g,z)\times S(x,z)$. The full seismogram is the superposition of all these arrivals. Note that this equation produces a complex valued dataset as output. The recorded seismic data is the real part of the complex dataset.

The migration and migration/inversion equations in chapter ?? can be similarly transformed to obtain the following equations.

Migration:

$$\hat{S}(x,z) = \int\int A(s,x,z)A(x,g,z)\, e^{-i[\phi(s,x,z)+\phi(x,g,z)]}P(g,s,\tau(s,x,z)+\tau(x,g,z))\ ds\ dg$$

Migration/inversion:

$$\hat{S}(x,z) = \frac{\int\int A(s,x,z)A(x,g,z)\, e^{-i[\phi(s,x,z)+\phi(x,g,z)]}P(g,s,\tau(s,x,z)+\tau(x,g,z))\ ds\ dg}{\int\int[A(s,x,z)A(x,g,z)]^2 ds\ dg}$$

These equations are applied to the real-valued trace P(g,s,t). The complex phase shift is implemented using the Hilbert transform:

$$e^{i\phi}P(t) = \cos(\phi)P(t) + \sin(\phi)\mathrm{H}(\mathrm{P}(\mathrm{T}))$$

where, H() denotes the Hilbert transform.

These imaging algorithms, and the band-limited Green's function algorithm described in the previous chapters, were implemented on a massively parallel computer, the Thinking Machines' CM-5.

1.2 Imaging the Marmousi dataset

The Marmousi model is a realistic model of the subsurface based on the geology of a field in offshore Angola. It has become a popular testbed for migration algorithms. Many groups have reported that Kirchhoff migration using first arriving traveltimes is unsuccessful in imaging the complex areas of the dataset (e.g. Geoltrain and Brac (?)). This is true even when the exact velocity field is used to calculate the traveltimes. The main reasons proposed for this include:

- The first arrivals may contain very little energy. Ideally the Kirchhoff algorithm should integrate over all arrivals at each location. If it only integrates one arrival and that arrival has little or no energy the imaging will not be successful.

- In order for ray tracing schemes and finite difference eikonal schemes to work the velocity field is often smoothed. Traveltimes and amplitudes in the smooth velocity field may not match wave propagation in the true velocity field.

Some asymptotic methods calculate more than one traveltime, and these methods can be used to generate maximum-amplitude traveltimes or multiple traveltime fields. However, even these methods do not image complex structures as well as wave equation based methods. One possible cause for this is that the Green's function may not be well represented by one (or more) traveltime picks. Another possible cause is that these traveltimes are calculated in the high-frequency limit, and are not representative of traveltimes in the seismic frequency band. In this chapter I will show that the second explanation is more likely in the Marmousi model. I show that a Kirchhoff algorithm that uses a single traveltime pick can generate a high quality image of the Marmousi dataset, as long as that traveltime is calculated in the seismic frequency band.

1.2.1 The Marmousi synthetic dataset

Figure 1.1 shows several shot records from the synthetic dataset. The data was modeled using a high order finite difference acoustic modeling algorithm. The shots have obvious non-hyperbolic moveout and back-scattered energy. This is an indication that prestack depth-migration may be necessary to image the data correctly. Figure 1.2 shows a near offset section from the dataset. The input data contains refractions and multiple reflections, but it does not contain any noise.

Since the dataset is the result of acoustic modeling it does not represent the true response of an *elastic* earth; it will not contain shear wave conversions or correct (elastic) amplitude versus offset behavior. If the P-waves and S-waves are strongly coupled, the solution of the acoustic wave equation is not a good approximation to the P-wave part of the solution of the elastic wave equation. In a real elastic medium there would be errors due to my use of the acoustic wave equation to approximate the Green's functions. Since this synthetic data was created using the acoustic wave equation, it will not test the validity of my use of the acoustic wave equation.

The dataset was modeled using a source that was approximately a 10-55Hz. minimum phase wavelet. The Kirchhoff migration algorithm that I use expects zero phase data as input. I therefore applied a minimum to zero phase conversion filter to the data before migration. After this filter it was apparent that the traveltimes of the water bottom reflection in the raw data did not match the expected traveltimes computed from the model. This is probably due to the effects of finite source and receiver depth and the particular implementation of the source function in the modeling scheme. To correct for this a 16ms. static shift was applied to the data to align

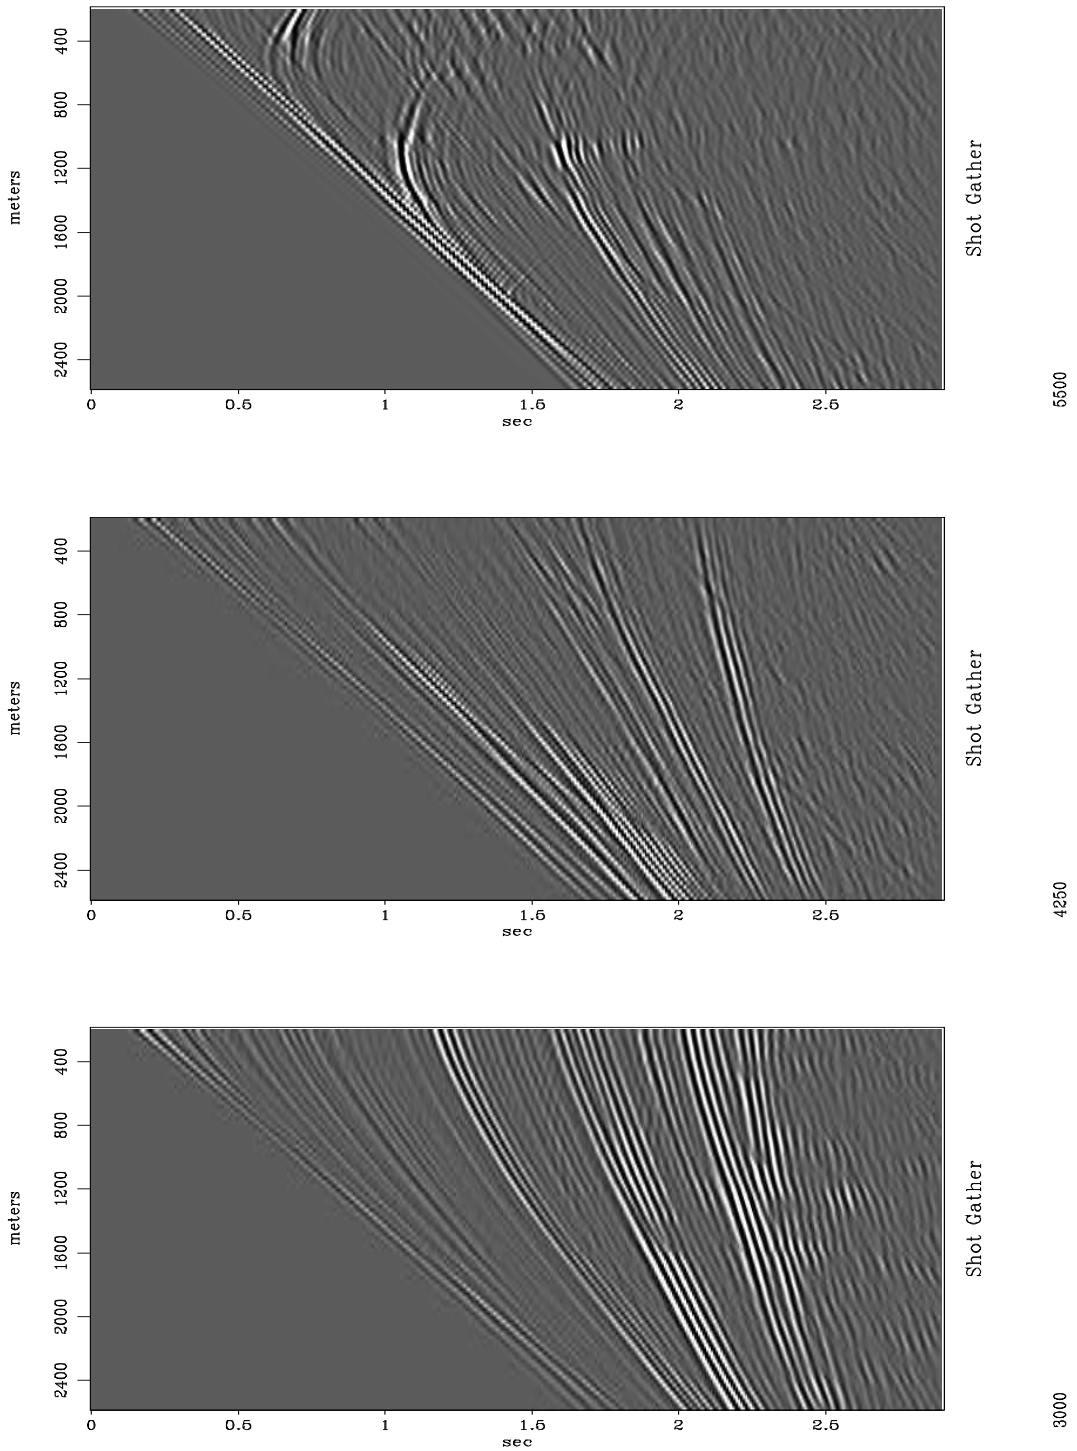

Figure 1.1: Marmousi synthetic data, shot gathers. MigMarm-marmshots [ER]

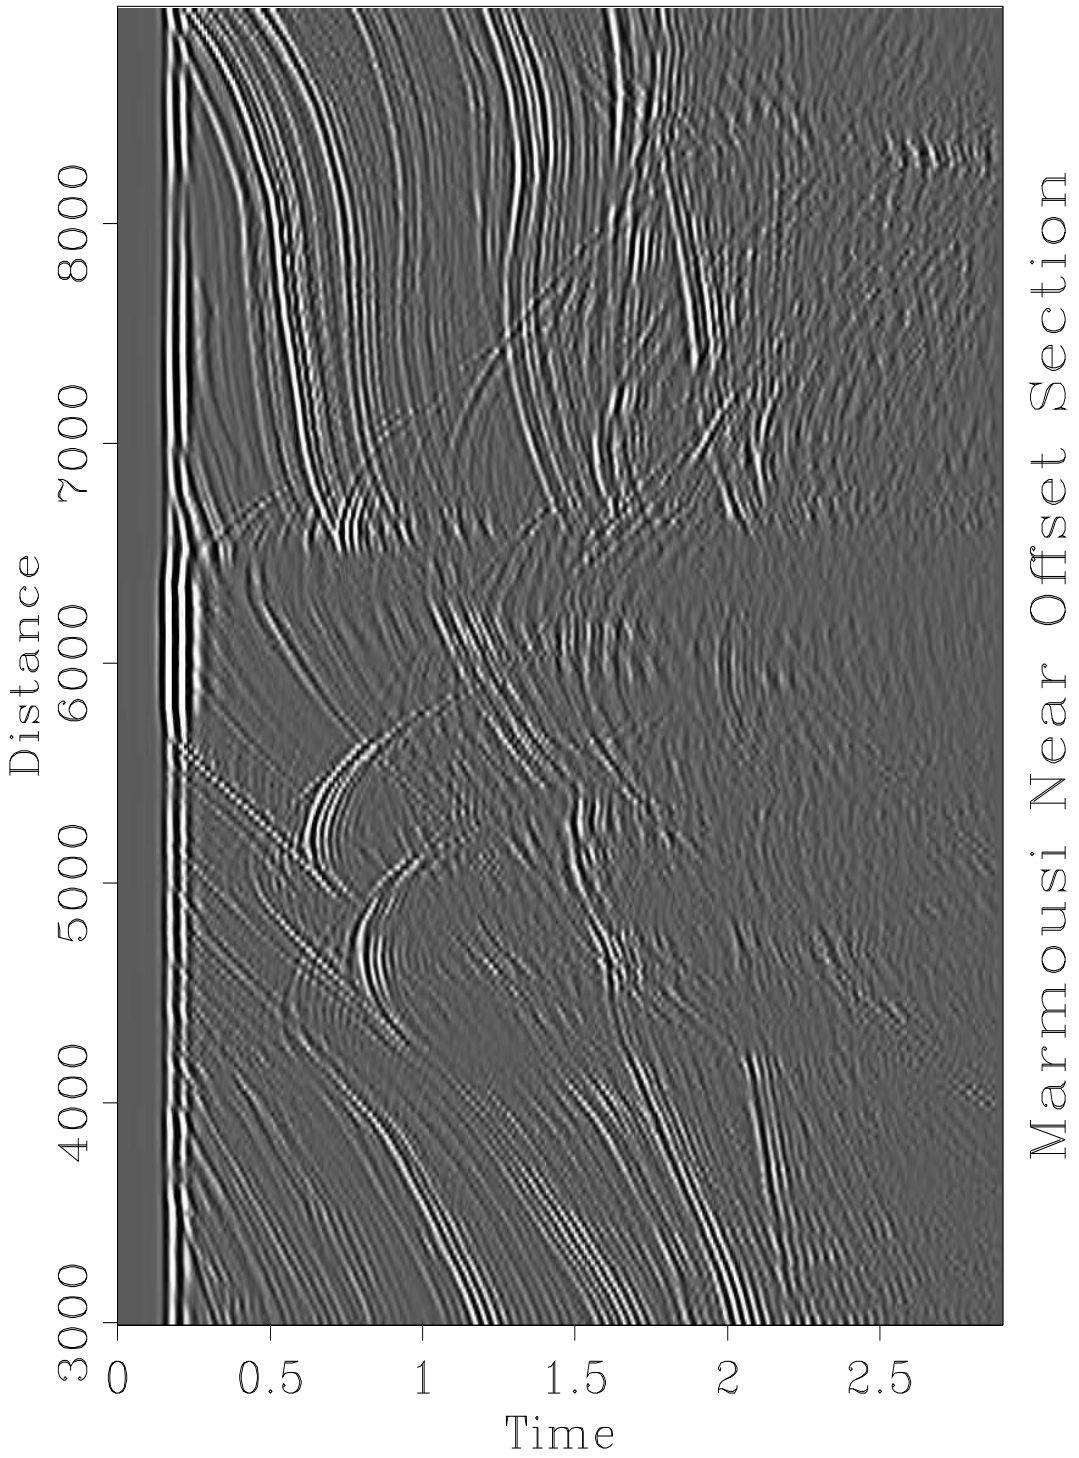

Figure 1.2: Marmousi synthetic data, near trace section. MigMarm-marmnear [ER]

the central lobe of the first arrival with the expected water bottom traveltimes.

The only other processing that was applied to the input data was a refraction mute. All events with an apparent velocity less than 1440m/s were muted from the input. In Figure 1.3 I compare one shot record before and after the pre-processing. No other preprocessing was applied to the input data. There was no attempt at deconvolution or multiple removal. In some of the images later in this chapter there are reverberations that may be due to multiples in the input data. I did not attempt to remove these from the input since the aim of this chapter is to compare Green's function estimation methods and not the preprocessing or the migration algorithm itself. These reverberations are indicated on the raw near-offset section in Figure 1.2.

Figure 1.4 shows the true Marmousi slowness model and Figure 1.5 shows a synthetic reflectivity section. The slowness and density models were combined to produce an impedance model and the vertical reflectivity at each location was calculated. The resulting data was convolved with a bandpass wavelet to create a model of the "ideal" imaging result. This is not exactly the response we would expect to see from a wave scattering from an interface in the subsurface, but it is a reasonable approximation to the desired output of the Kirchhoff migration/inversion algorithm. If amplitudes are not available we would only expect to see an equivalent structural image, and not the same relative amplitudes on the reflections.

1.2.2 Migration using standard algorithms

Figure 1.6 shows the result of using a commercial finite-difference shot-profile migration[1] to image the Marmousi dataset. This program uses explicit finite-difference extrapolators to downward continue the source wavefield and the received wavefield for every shot. At each depth an image is formed using the "correlation imaging condition". This imaging condition is the equivalent of using the adjoint operation in Kirchhoff imaging. This result can be regarded as a benchmark for structural images. Since it only attempts to recover the structure and not the true reflectivity, it should not be compared to the results of migration/inversion. The problems with shot profile migration are two-fold. First it is an expensive process because it extrapolates every frequency to all depths. Second, it requires regular and fine spacing of the receivers in a shot. While this is common in 2-D, it is very uncommon in 3-D data. The main advantages of Kirchhoff methods are that they can handle sparse and irregular sampling much more easily than downwards extrapolation methods, and that they are cheaper than full-wavefield finite-difference methods.

All of the Kirchhoff migration results shown in this chapter were created using the same program. This program is a very simple implementation of equations (1.1) and (1.1). Simple linear interpolation is used to extract input data at non-integral time locations. It is not a "commercial quality" program, and no attempt has been made to implement features that can create a more pleasing image. In particular, no attempt is made to handle aliasing problems caused by very steep summation

[1]The shot-profile migration in the ProMAX seismic processing system from Advance Geophysical

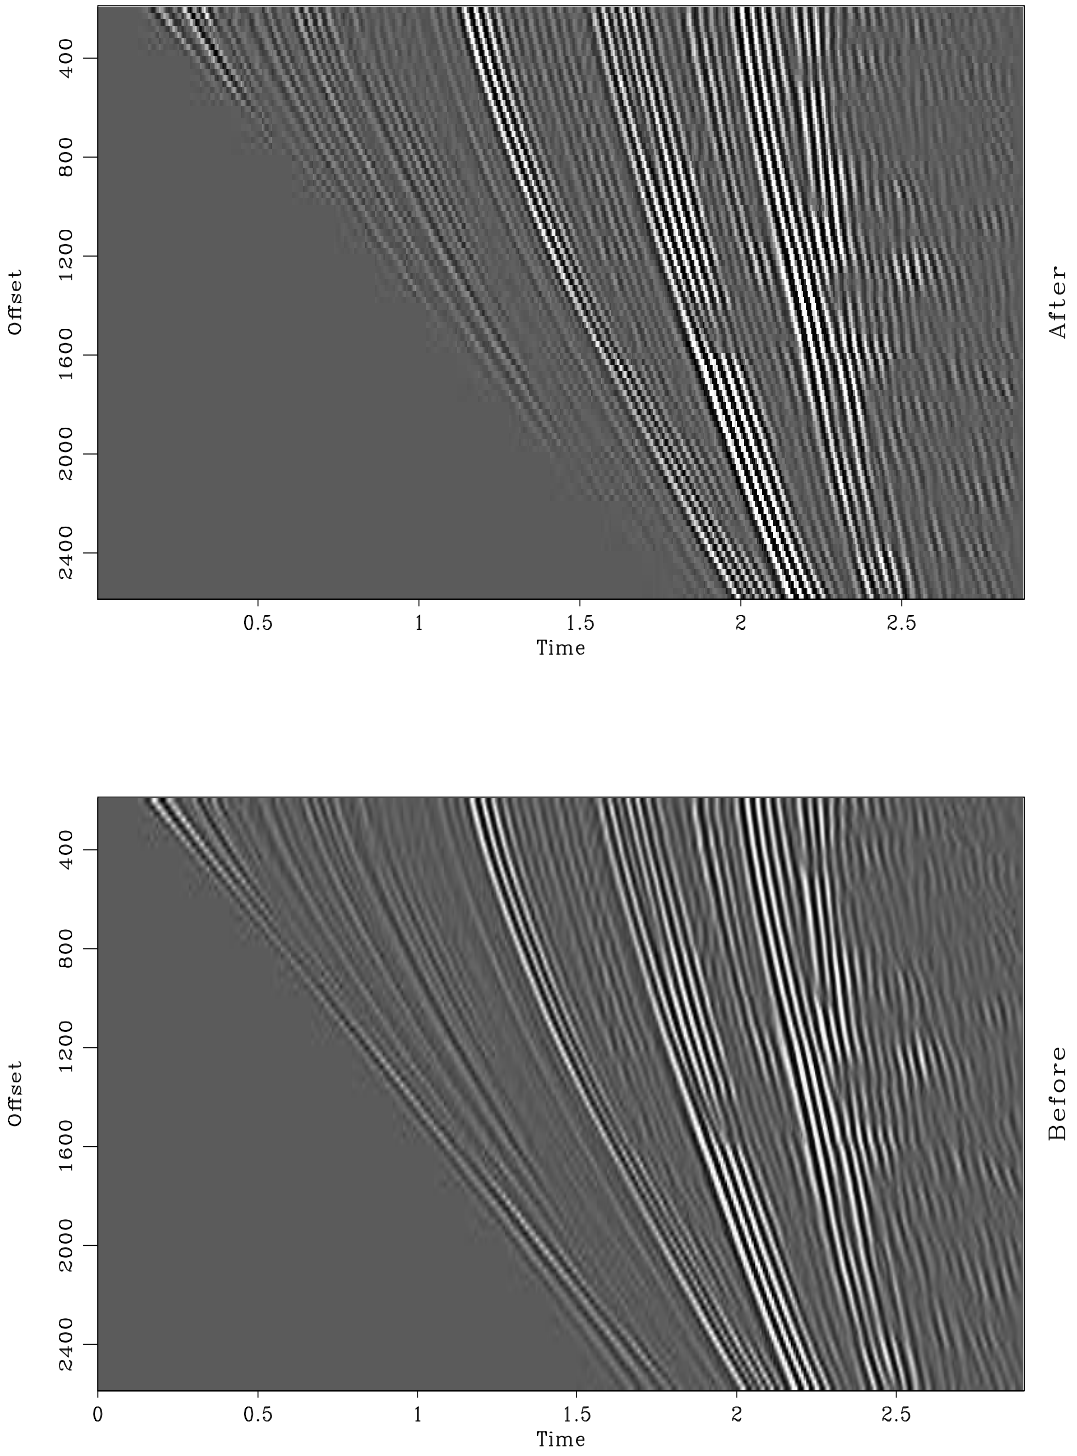

Figure 1.3: Marmousi shot record before and after pre-processing: De-phase, static shift and mute. MigMarm-shot-prep [ER]

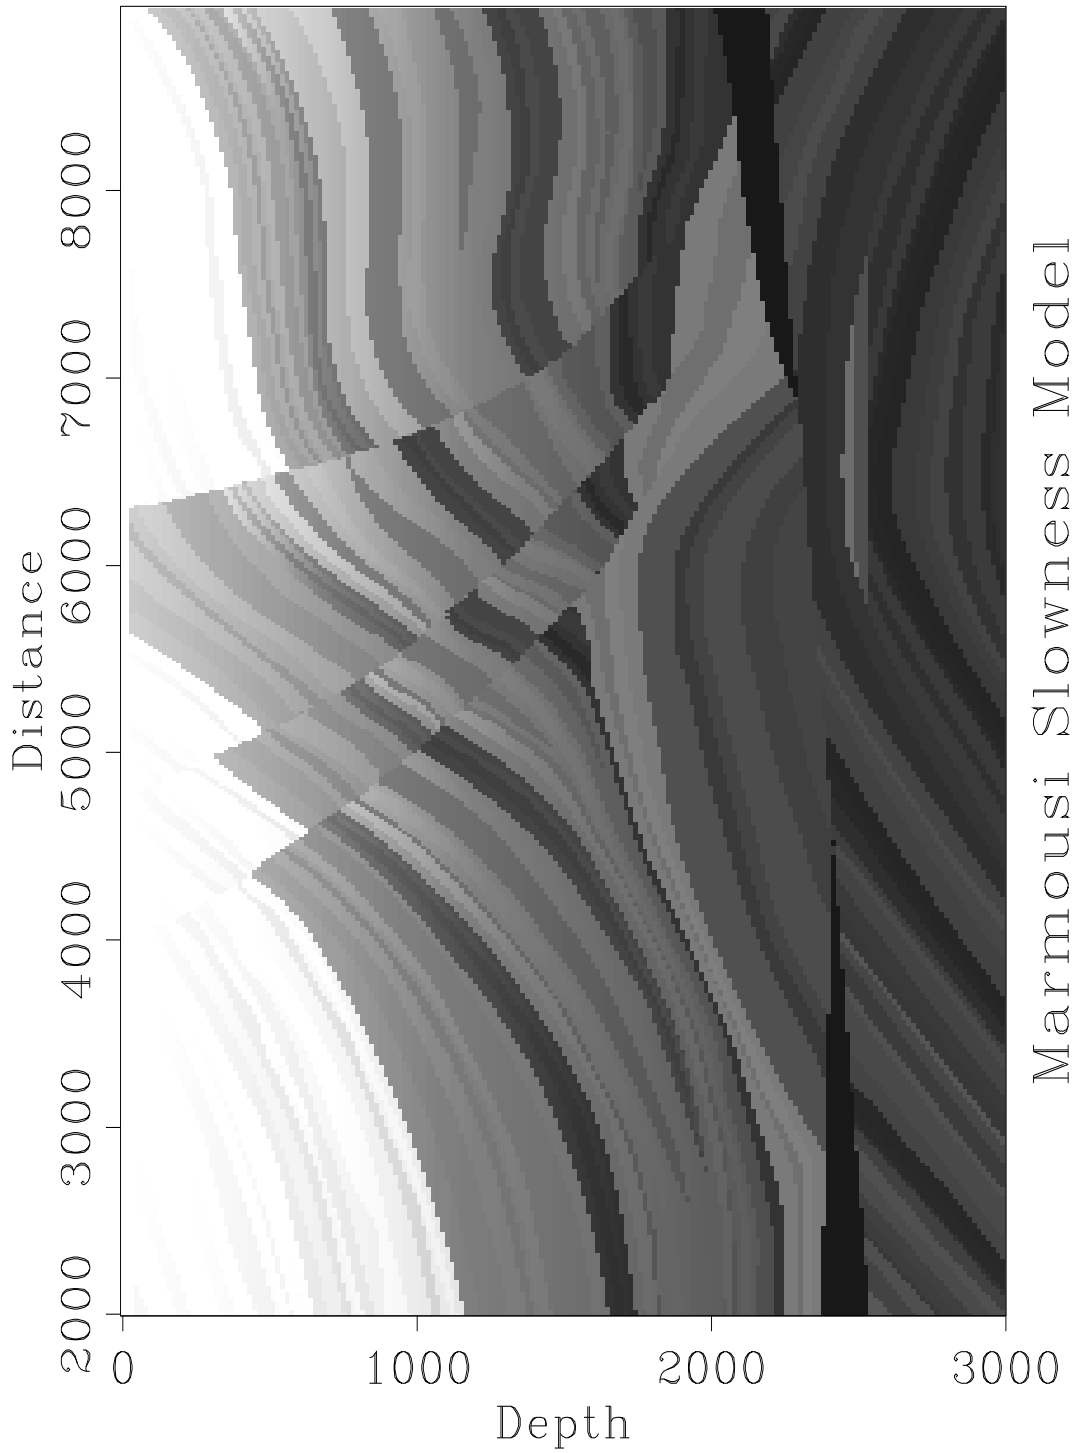

Figure 1.4: Marmousi true slowness model. MigMarm-marmslow [ER]

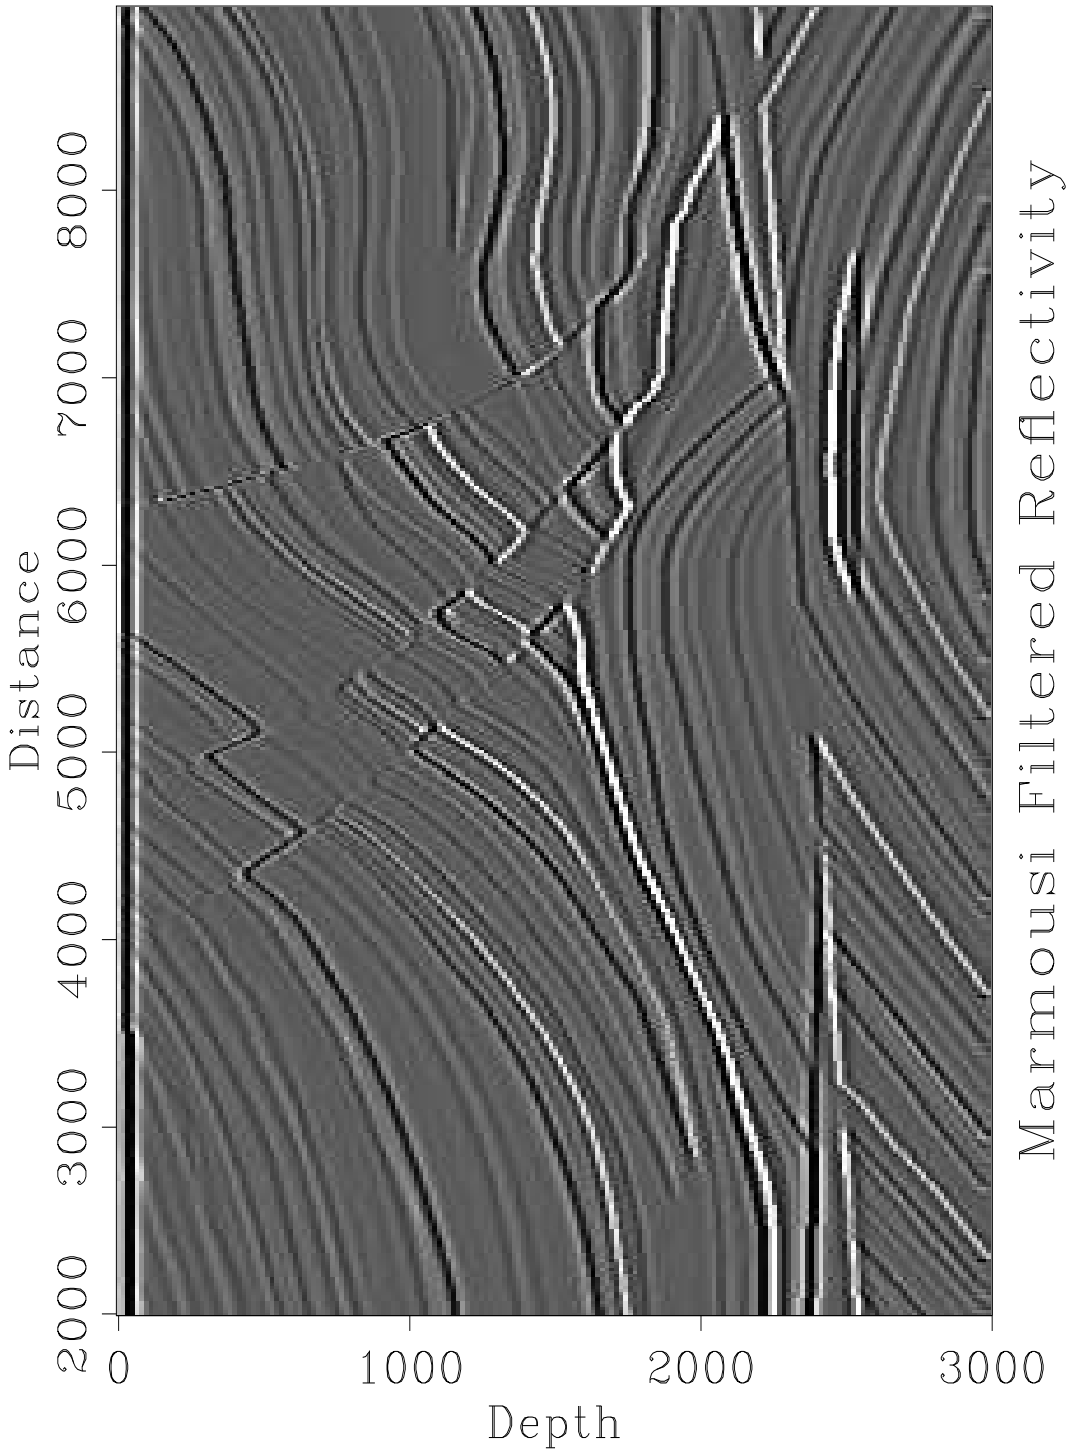

Figure 1.5: Synthetic reflectivity section for the Marmousi model. [ER] MigMarm-marmrefl

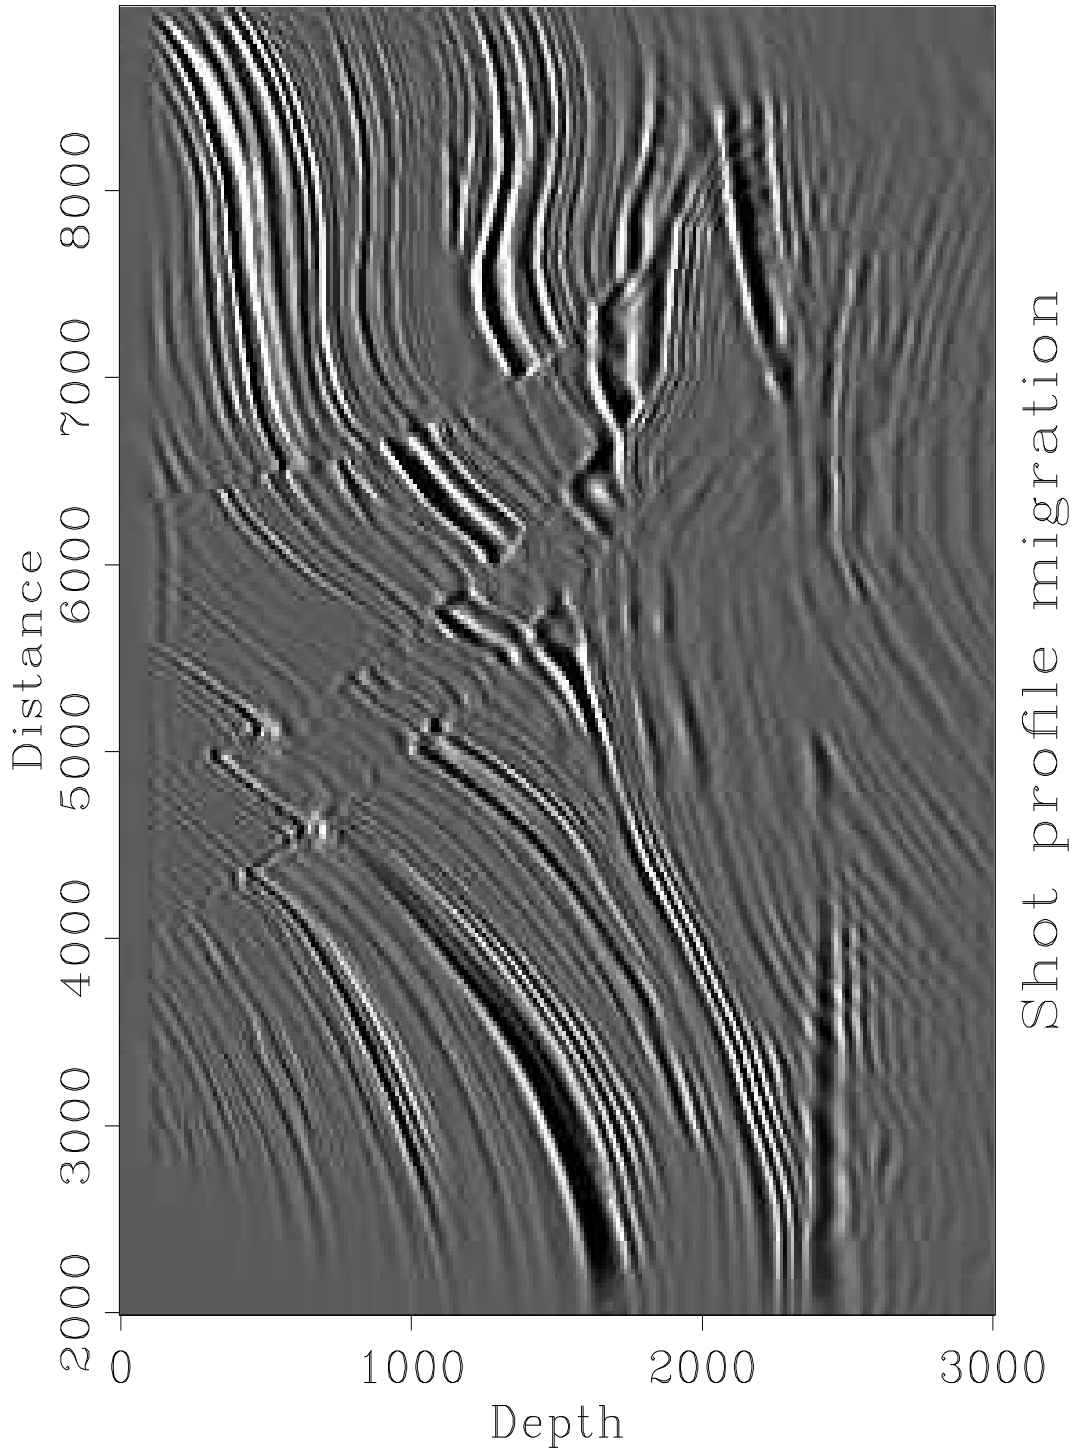

Figure 1.6: Marmousi dataset migrated using a commercial shot profile migration program. This program performs downwards extrapolation of source and receiver wavefields by a finite-difference algorithm. MigMarm-mig-shot-prof [CR]

trajectories. The effects of aliasing are not noticeable in the stacked images, but they are more evident in the CDP gathers shown later in this chapter (Figures 1.14–1.16). The migration output could also be improved by applying a post-migration mute to the CDP gathers, or by making use of signal to noise ratio estimates to weight the stack. I did not attempt to implement these features, as the aim of this chapter is to compare the effects of different methods for computing the Green's functions. What is important is that exactly the same migration program, whatever its flaws, is used for all the different methods.

The Kirchhoff migration program has the ability to handle different parametric approximations to the Green's function. The Green's functions can be provided as a traveltime map only; traveltime and phase maps; or traveltime, amplitude, and phase maps. If phase maps are not available the phase is assumed to be zero. If amplitude maps are not available the amplitude is assumed to be one. When amplitudes are available, there are two options for their application. The operator can either be applied as an "adjoint-modeling" operation described by Equation (1.1) or as the diagonalized "migration/inversion" operator described by equation (1.1).

Figure 1.7 shows the result of a migration using first arrival traveltimes calculated using an upwind finite-difference solution to the eikonal equation. While much of the upper portion of the model has been well imaged, the target zone below 2000m is almost totally lost. This method has performed well in the simple parts of the velocity model but it has failed in the complex regions. This is the type of result that has caused renewed interest in other methods of calculating the traveltimes.

Figure 1.8 shows the result of migration using just the traveltime and phase maps for maximum amplitude arrivals calculated by paraxial ray tracing. This is a much better image of the subsurface, and almost all features of interest have been imaged. The top of the trap structure at coordinates (6500,2500) is visible, but the flat spot below it is not clear. There are a few regions that have not been well imaged. The blocky structures at (6000,1700) and the domed structure below them at (5000,2000) are unclear. However, in general this is a very good result. It shows the advantage of using maximum amplitude arrivals rather than first arrivals.

Figure 1.9 shows the image created by migration/inversion using the traveltime, phase, and amplitude information from paraxial ray tracing. One obvious problem is the loss of amplitude at the top of the section. This is seen on most of the paraxial ray tracing migration/inversion results. It is probably caused by a failure to account for free-surface effects in the Green's functions. In the band-limited traveltime estimation a $\cos\theta$ amplitude function is applied at the source and receiver to approximate the near surface effects. The lack of this term in the paraxial ray tracing would account for the higher amplitudes at the bottom of the section, since the expected amplitude decay in the input data may not be present. Apart from this loss of amplitude at the top of the section the image is very good. The correct relative amplitudes are seen on the different reflectors. Using the amplitudes in the migration has improved the imaging a little, the trap is clearer and some of the blocky structure is better imaged.

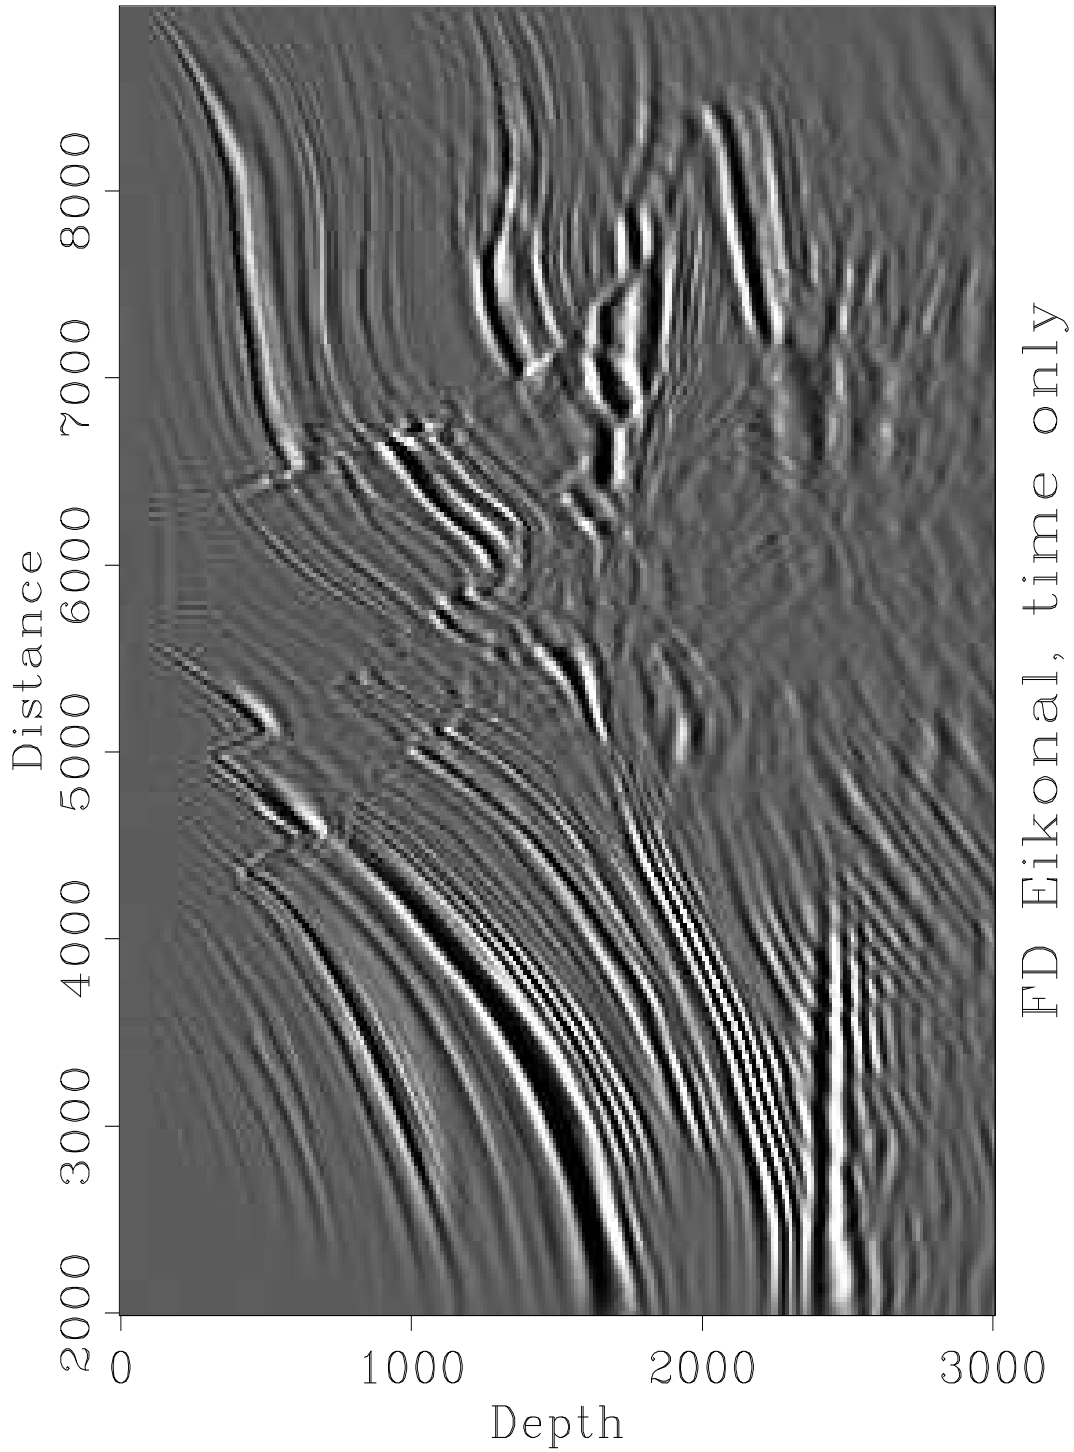

Figure 1.7: Marmousi dataset migrated using traveltimes from finite difference solution to the eikonal equation. MigMarm-eikonal-t [CR]

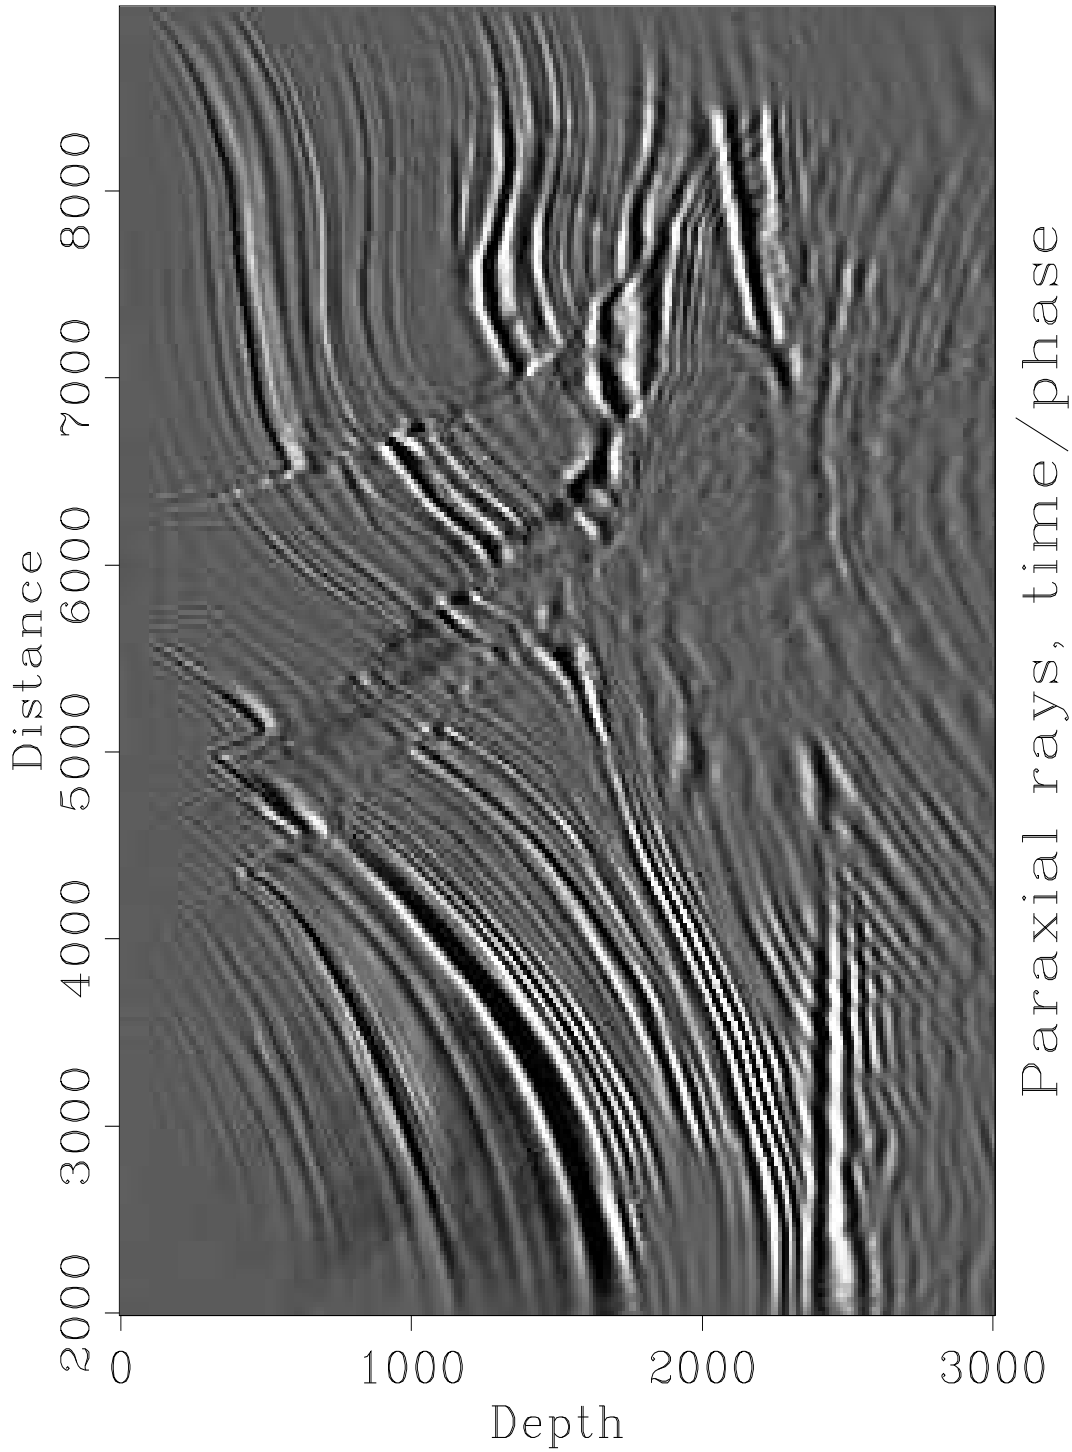

Figure 1.8: Marmousi dataset migrated using traveltime and phase maps calculated by paraxial ray tracing. MigMarm-paraxial-tp [NR]

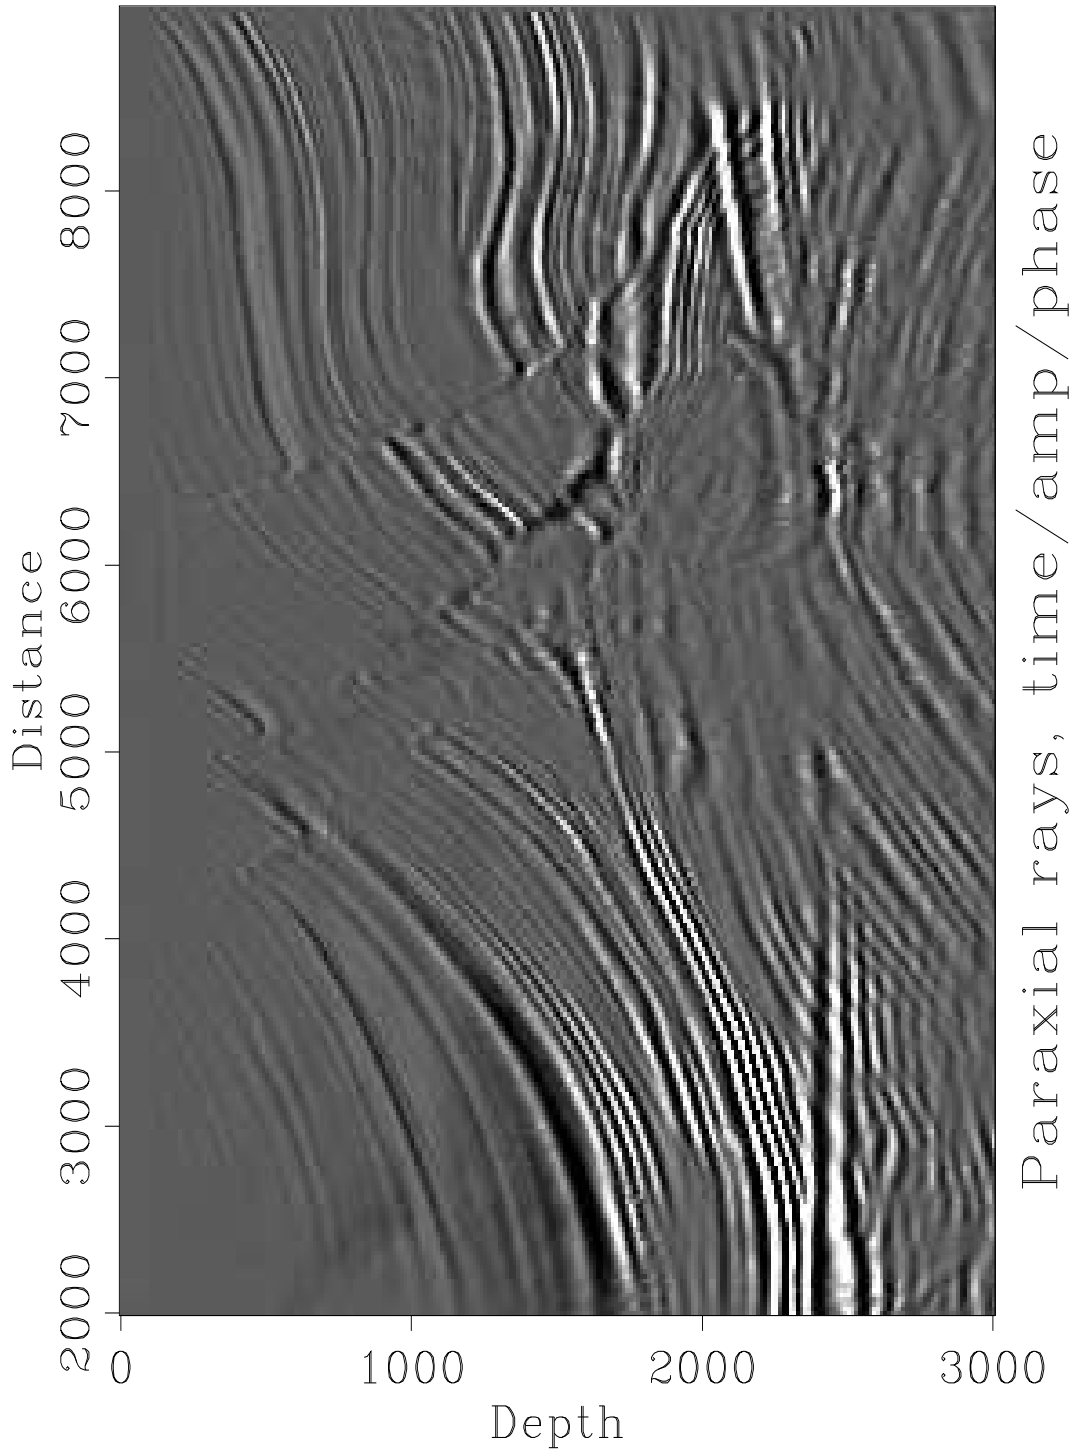

Figure 1.9: Marmousi dataset migrated using traveltime, amplitude, and phase maps calculated by paraxial ray tracing. MigMarm-paraxial-tap [NR]

1.2.3 Kirchhoff migration using band-limited Green's functions

Figure 1.10 shows the result of migration using the traveltimes and phases from the band-limited Green's functions. To create these traveltimes I used sixteen frequencies in the range 10-60Hz. for the Green's function estimation. This image is even better than the image created using paraxial ray tracing. The blocky structures and the domed section that were missing in the previous image have been recovered and the target trap is well imaged. Even the flat spot at the base of the reservoir is visible.

Figure 1.11 was created by using the traveltime, amplitude, and phase maps in an adjoint-modeling scheme. This result should be compared with Figure 1.6, the output of the finite-difference shot-profile migration. Both schemes are designed to produce a structural image rather than a true reflectivity estimate. By applying the adjoint operator the signal to noise ratio in the structural image is maximized. The shot-profile migration performs two depth extrapolations of every frequency at each shot location (one to downward continue the received wavefield and one to model the source wavefield). The band-limited Green's function computation extrapolates 16 frequencies at each surface location, then performs a Kirchhoff migration. In most reasonable geometries this will be much cheaper.

Figure 1.12 shows the result of migration/inversion using the amplitude information. This shows good recovery of the true relative reflectivities and a slightly crisper image than using traveltime and phase alone. In Figure 1.13 a similar result has been obtained using only eight frequencies in the range 10-35Hz. This was the smallest number of frequencies that I was able to use in the traveltime estimation, while still producing a good image. With only eight frequencies to be propagated the cost of producing the Green's function tables is only a few times that of the finite difference solution to the eikonal equation but the image is clearly superior!

1.2.4 Comparison of CDP gathers

As I mentioned in the introduction to this thesis, the imaging and velocity estimation problems are closely coupled. One common way to evaluate the correctness of the velocity model is to look at CDP gathers after common offset or common shot migrations. If the correct velocity model has been used the migrated image should be in the same position on all offsets. There may be variation in amplitude due to angular variation of reflectivity but the images should all line up. If the velocity model is wrong it will affect the different offsets in an offset dependent manner. The resulting images will not be aligned on the different offset panels. Many velocity estimation methods use the residual moveout of events within CDP gathers as a misfit measure and use this information to update the velocity model (**?**; **?**).

Figures 1.14 to 1.16 show CDP gathers after migration. These are selections of the same CDP locations from the different constant offset migrations. If the images are consistent on the different offsets then the velocity model and Green's function

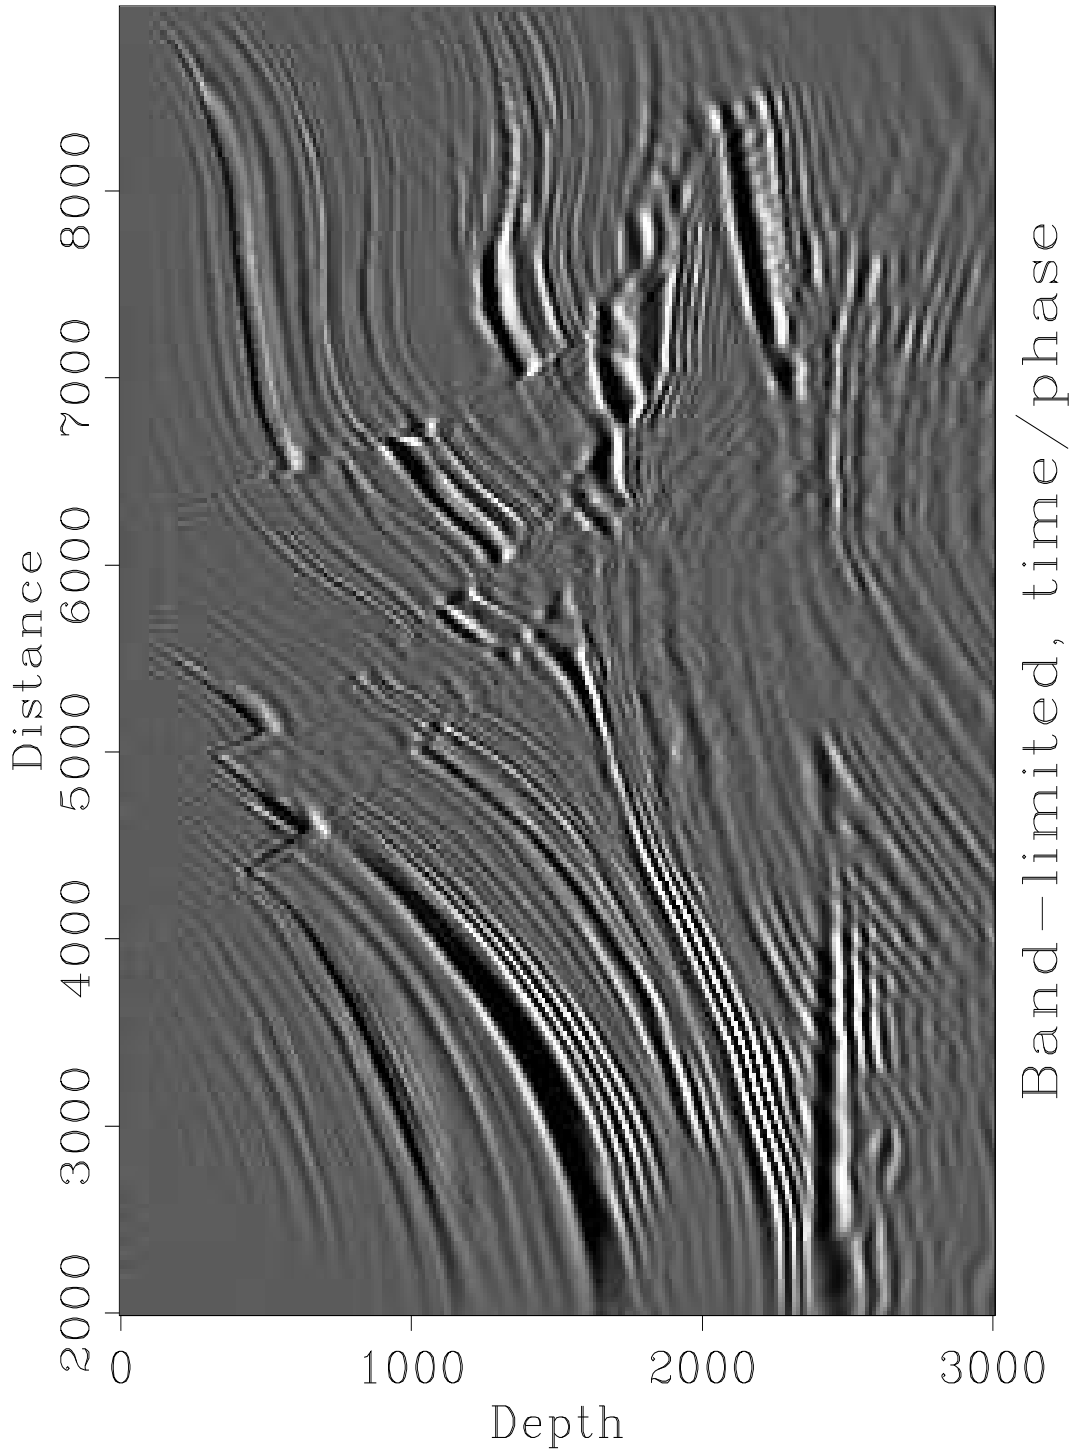

Figure 1.10: Marmousi dataset migrated using band-limited traveltimes and phase maps. MigMarm-marm-tp [CR]

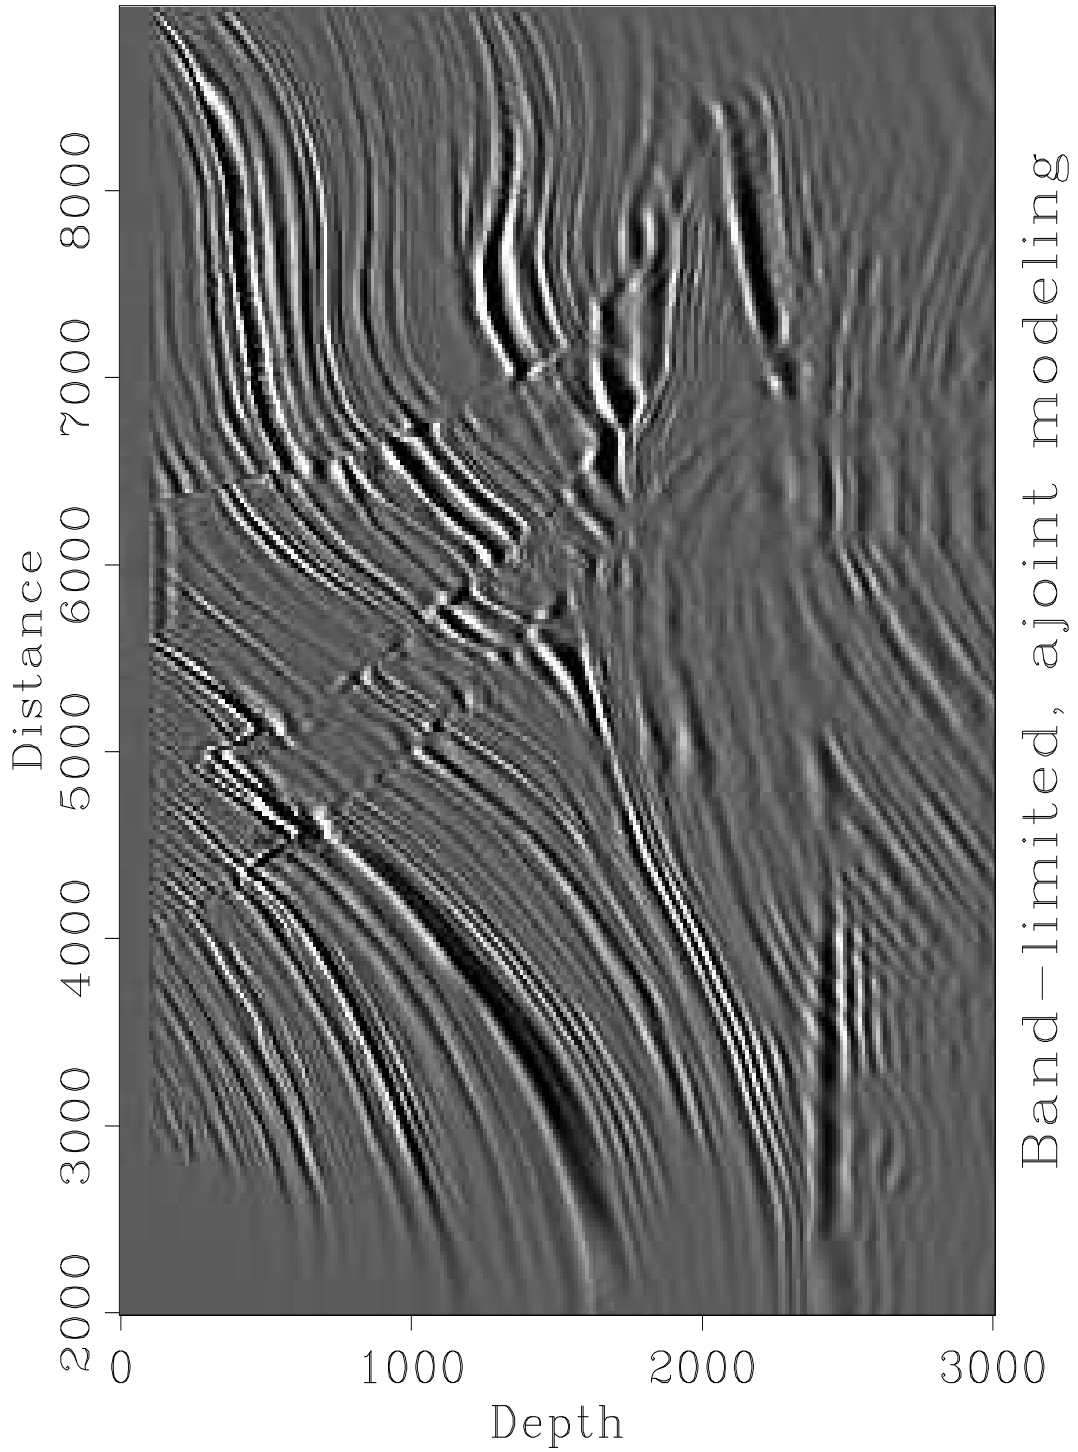

Figure 1.11: Marmousi dataset migrated using band-limited traveltime/amplitudes/phase maps in an adjoint modeling algorithm. Compare with Figure 1.6. MigMarm-marm-adj [CR]

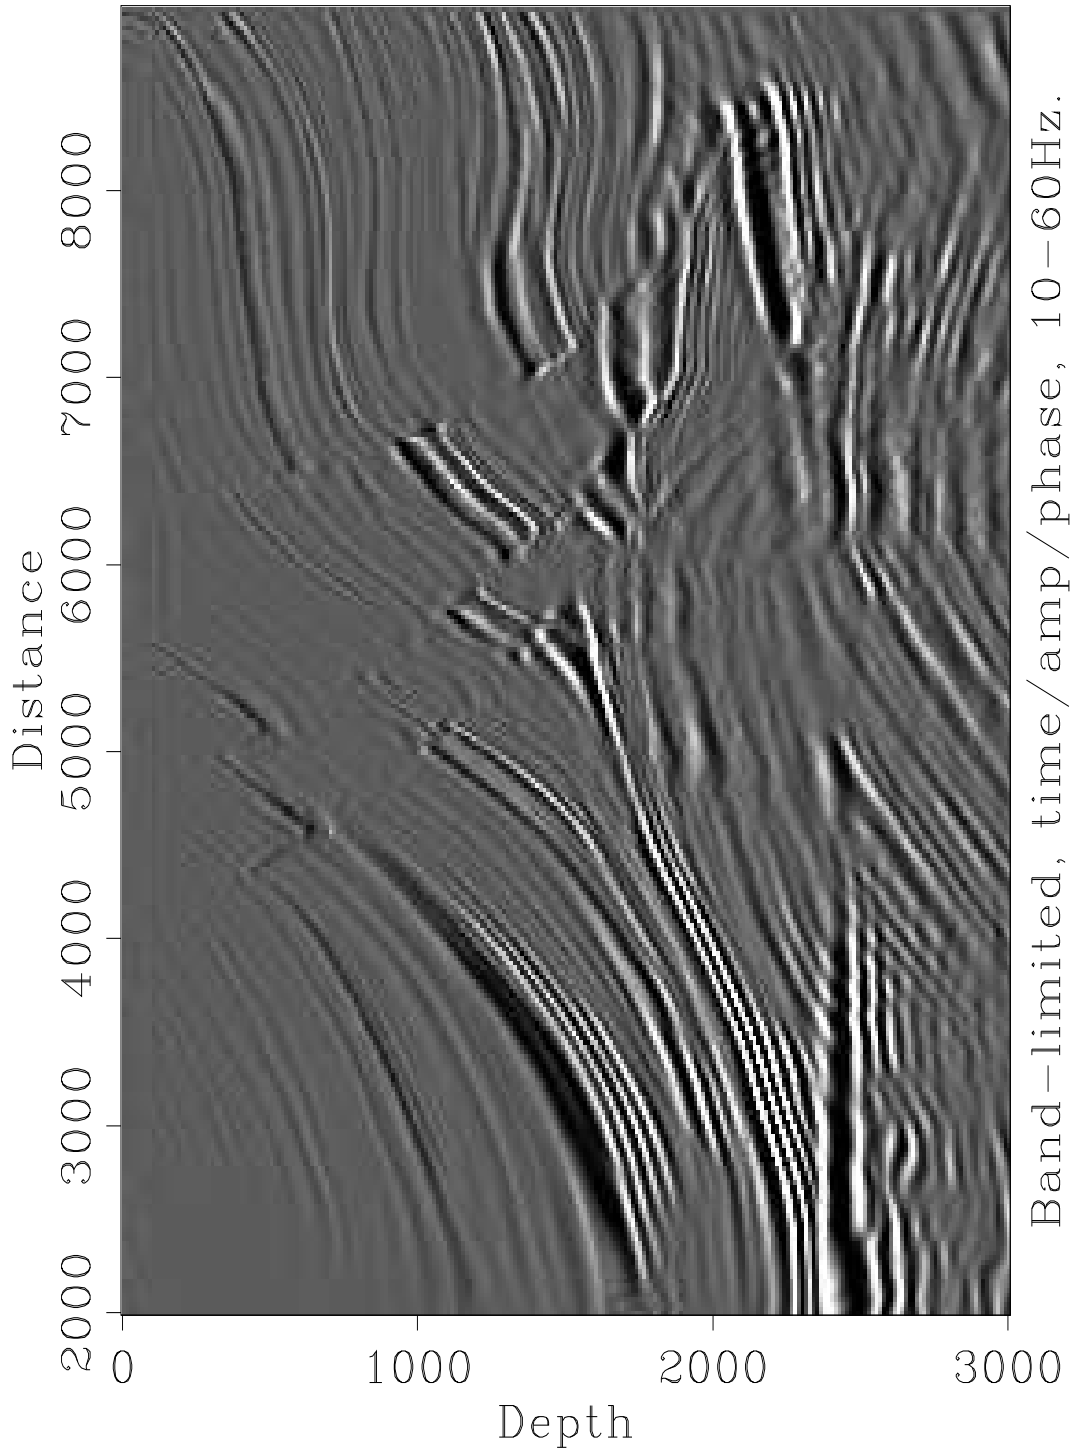

Figure 1.12: Marmousi dataset migration/inversion using band-limited traveltime, amplitude, and phase maps. Traveltimes were estimated from 16 frequencies in the 10-60Hz frequency band. Compare with Figure 1.9. MigMarm-marm-tap [CR]

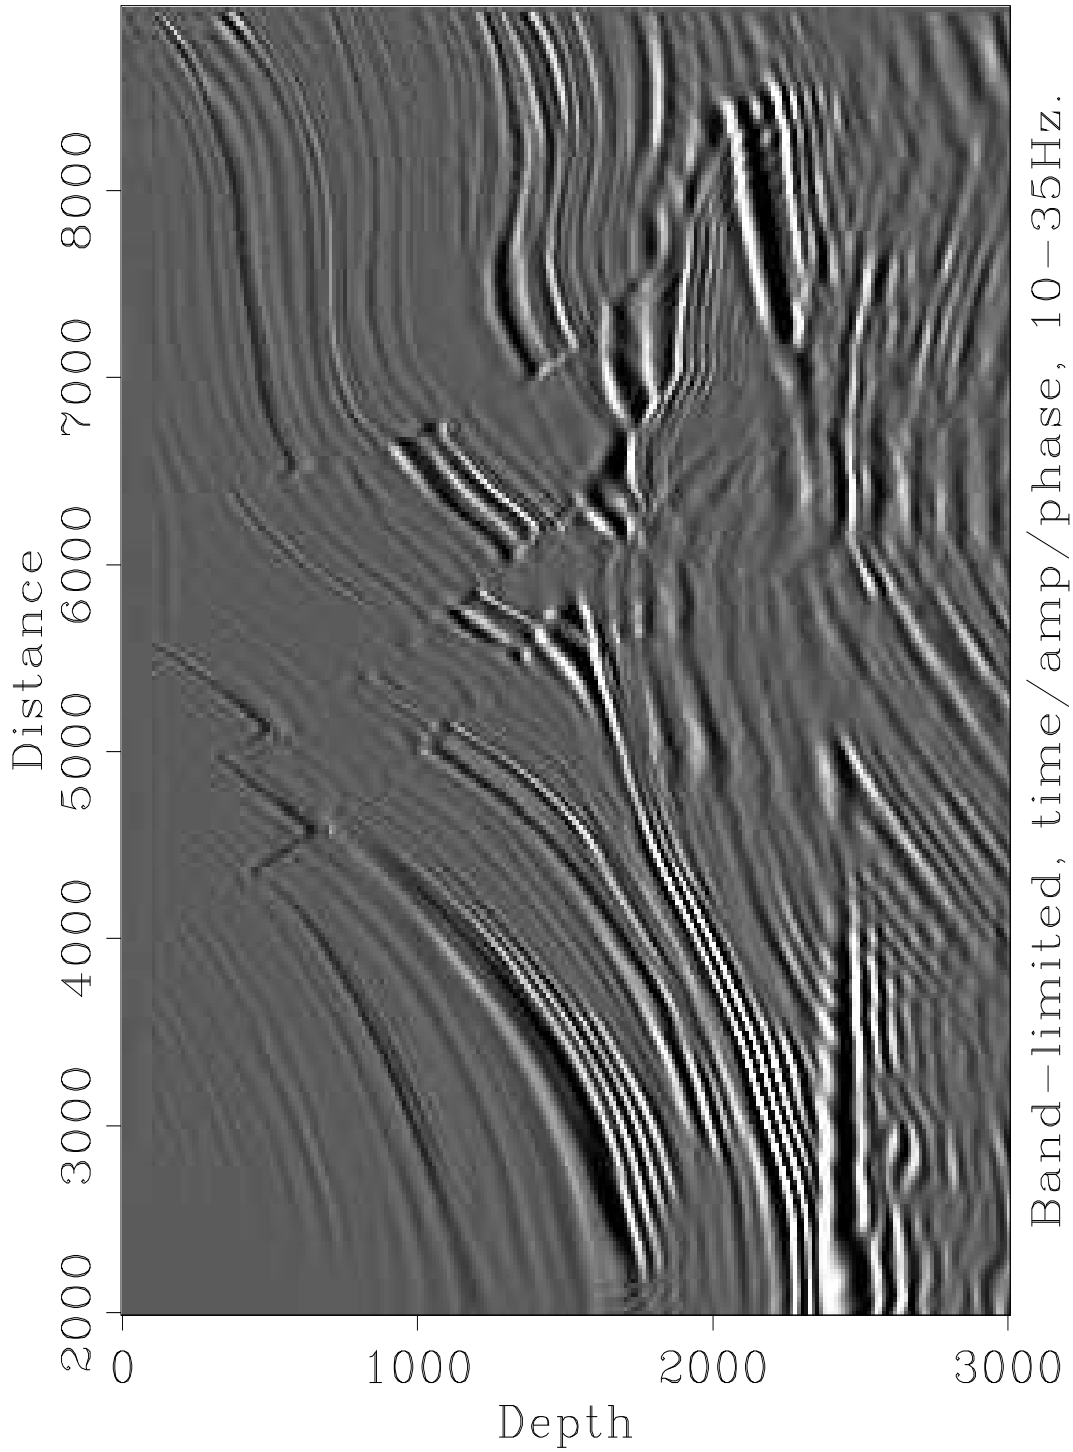

Figure 1.13: Marmousi dataset migrated using band-limited traveltime, amplitude, and phase maps. Traveltimes were estimated from 8 frequencies in the 10-35Hz frequency band.Compare with Figure 1.12. MigMarm-lomarm-tap [CR]

calculation have matched the input data well.

It is clear that at many locations on 1.14 the images are not consistent across different offsets. These images were created using the Green's functions calculated using a finite-difference solution to the eikonal equation. The events are not flat and they sometimes have phase changes with offset. This is very disappointing; if we were to use coherency of images across offset as a measure of the fitness of the velocity model, we would assume that this was a poor model. In fact, it is the perfect model! It is the traveltime estimation method that is unfit, not the model.

The paraxial ray tracing and the band-limited Green's functions give much more consistent results. The band-limited Green's function result is clearly better in those areas that it images well, particularly on the panel for CDP location 7000. Figure 1.17 shows a close up of CDP 7000 for all three methods. The band-limited Green's function result is the most coherent of the three panels. This is the result that we expect; when the true velocity model is used the images should be coherent across offset. If the CDP gathers are to be used for velocity model updating we need an additional result. When the wrong velocity model is used, the CDP gathers should be less consistent as a function of offset, this remains to be demonstrated.

There are still non-horizontal events on these panels. This is not surprising, since this migration only images one arrival at each location. The signal associated with the lower energy arrivals is still present and it will not be imaged consistently across all offsets. If the other arrivals are close in energy to the maximum-energy arrival we would expect significant degradation of the image. In these situations we would expect to have to use a Green's function estimation method, and Kirchhoff migration, that accounted for multiple arrivals. It is encouraging that such high quality images are created with only one arrival. It suggests that, even in a very complex model, a single event model may be a reasonable way to describe the Green's function.

1.2.5 Imaging using an approximate velocity model

It might be claimed that the previous images are unrealistic. We will never have a perfect velocity model with resolution as good as that shown in figure 1.4. In the previous section I showed that the band-limited Green's functions gave a better image when the true velocity model was used. In this section I demonstrate that they still give a better image when an approximate model is used.

Figure 1.18 shows the approximate model that I used. It was created by choosing significant structures in the true model and replacing the true slowness within each structure with a smoothed version of the true slowness. It is intended to be the type of model that a "perfect geologist and perfect geophysicist" might have created by using clever velocity estimation algorithms, a-priori knowledge of the region's geology, and a little luck. Figure 1.19 shows the image created by migration/inversion in the simplified model using band-limited Green's functions. While it is not as good as the image created using the true model it is still very clear. Figure 1.20 shows the result obtained using Green's function calculated by paraxial ray tracing in the simplified

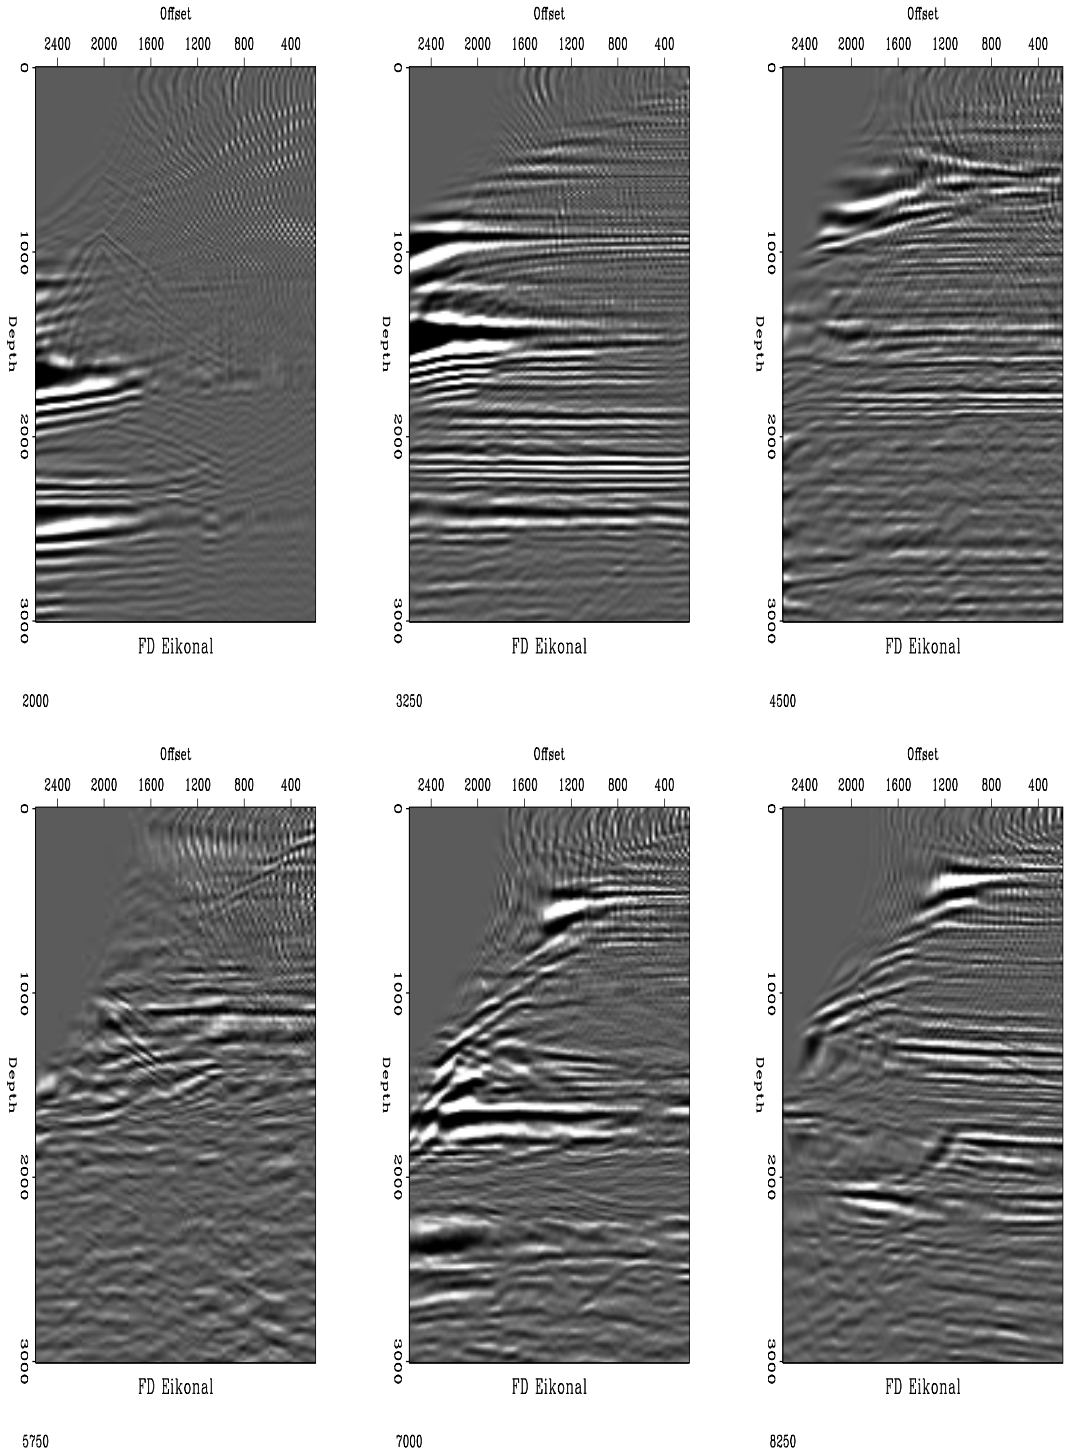

Figure 1.14: CDP gathers computed using first arriving traveltimes. [CR]

MigMarm-eikonal-cdp

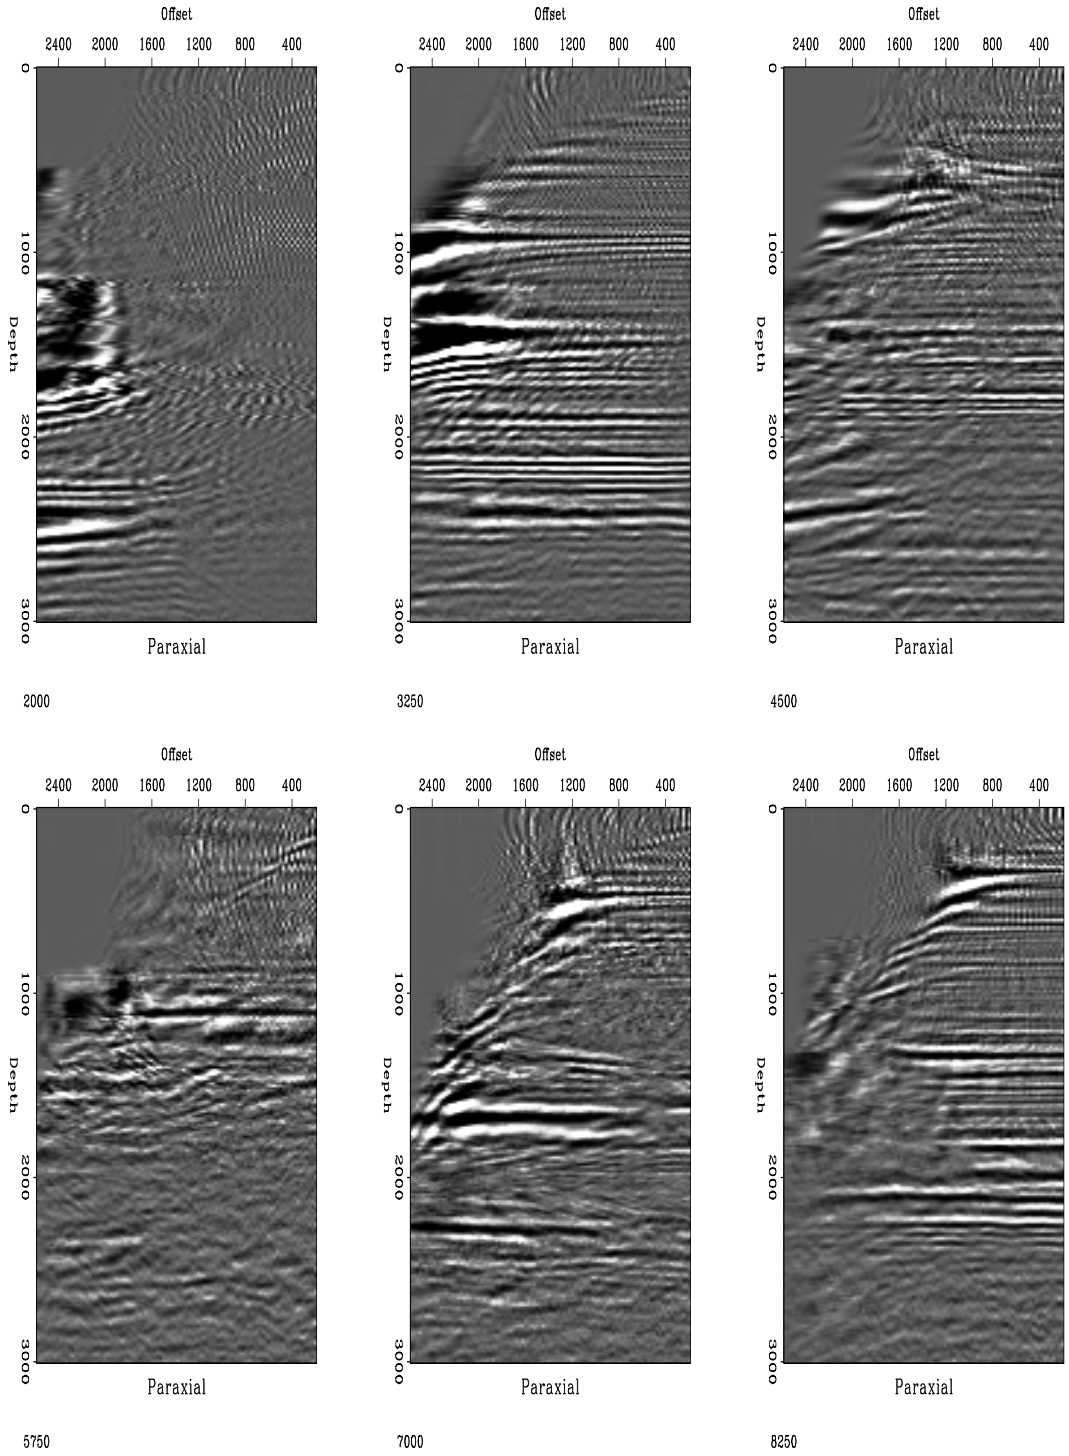

Figure 1.15: CDP gathers computed using traveltimes and phases from paraxial ray tracing. MigMarm-paraxial-cdp [NR]

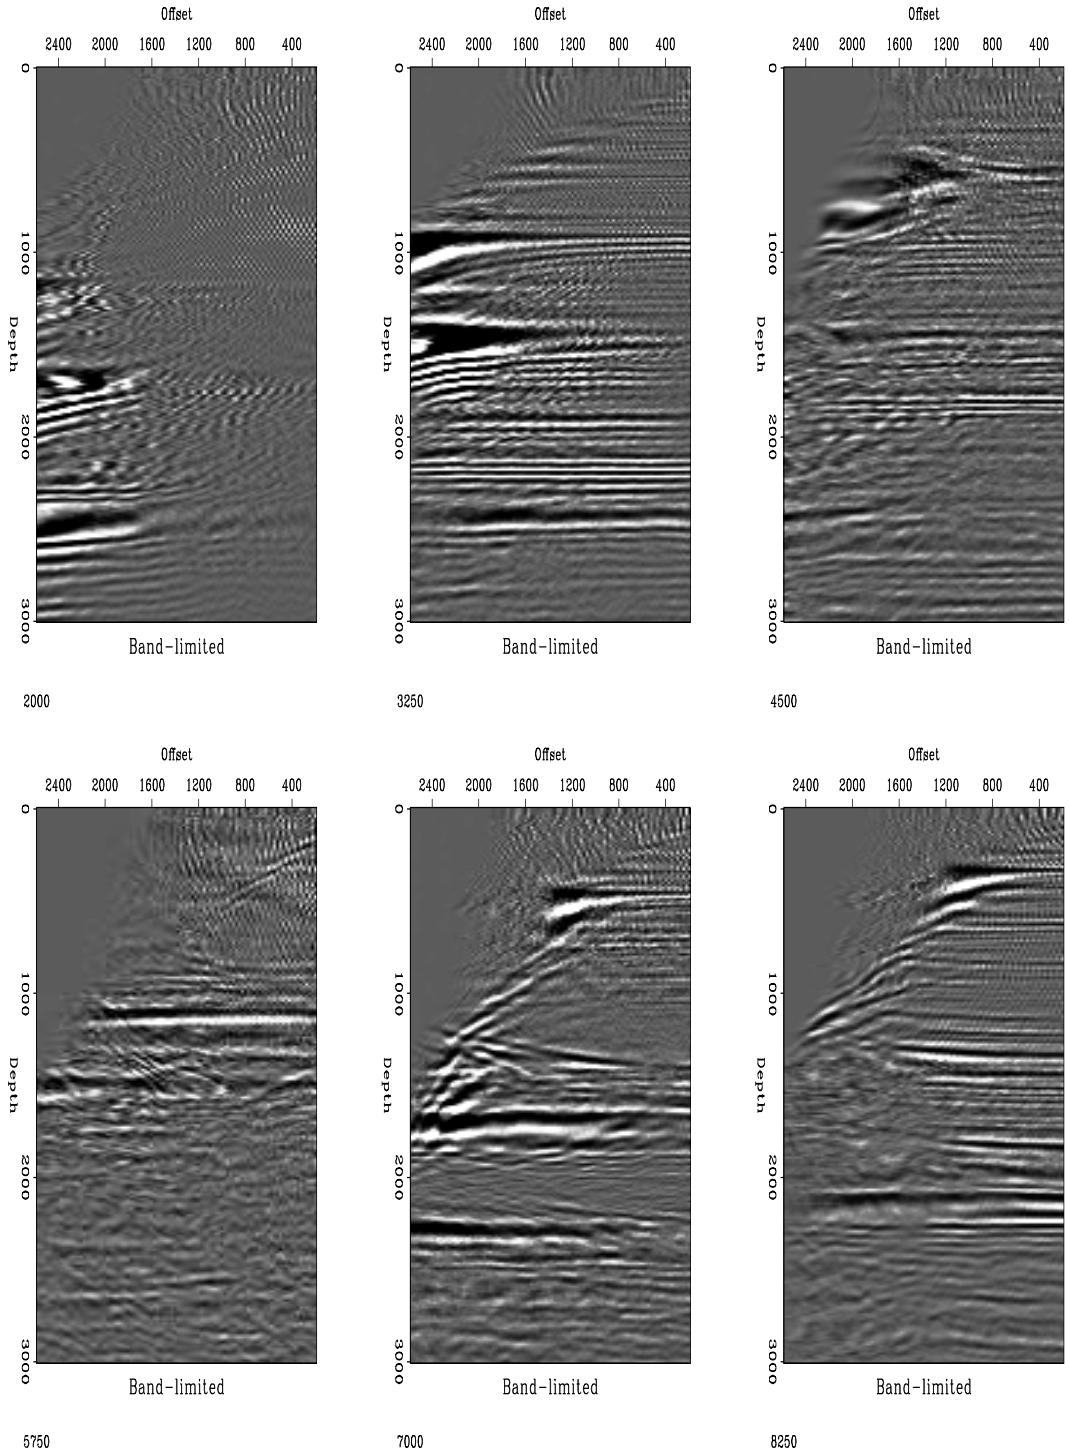

Figure 1.16: CDP gathers computed using traveltimes and phases from band-limited Green's functions. MigMarm-marm-cdp [CR]

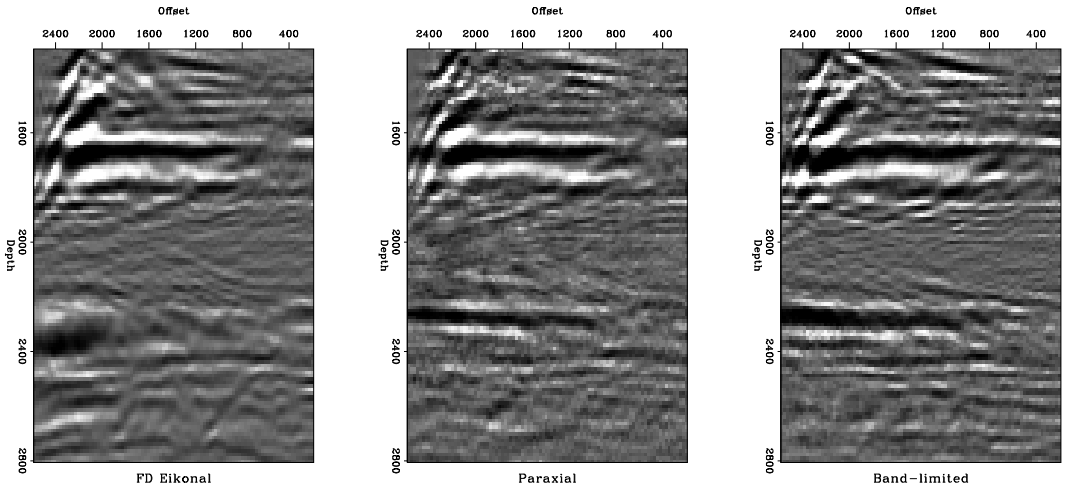

Figure 1.17: A close up of CDP 7000 for three methods. On the left migration using first arriving traveltimes. In the center migration using traveltimes from paraxial ray tracing. On the right migration using band-limited maximum-energy traveltimes. MigMarm-all-cdp [CR]

model. It is not as clear as the result obtained using band-limited Green's functions.

This comparison illustrates that the band-limited Green's functions are still an improvement on paraxial ray tracing, even when the model is simplified. In this particular case there are still large velocity discontinuities in the model. I have not demonstrated that band-limited Green's functions are better than paraxial ray tracing when a smooth model, one that is ray valid, is used. This remains a topic worthy of further study. What I would claim is this: when a velocity estimation procedure has produced a discontinuous model it is better to calculate Green's functions with a method that is valid in a discontinuous model, rather than use a method that requires the hard won velocity model to be artificially smoothed.

Figure 1.21 shows a comparison of CDP gather 7000 for the simplified model and the true model. Both results were obtained using band-limited Green's functions. The use of the simplified model has degraded the continuity as a function of offset. This is a pleasing result, since it suggests that coherency with offset is sensitive to changes in the velocity model. It also suggests that this particular simplified model is not the best possible model. We can obtain better continuity by using a model that is closer to the true model. The trick, as always, will be in finding a scheme for updating the model so that the image and the continuity within CDPs improves.

1.3 Summary

In this chapter I have shown that Kirchhoff migration can produce images comparable to full wavefield methods as long as a representative traveltime is chosen. In complex models a first-arrival traveltime field may not be a good choice, and it has been shown

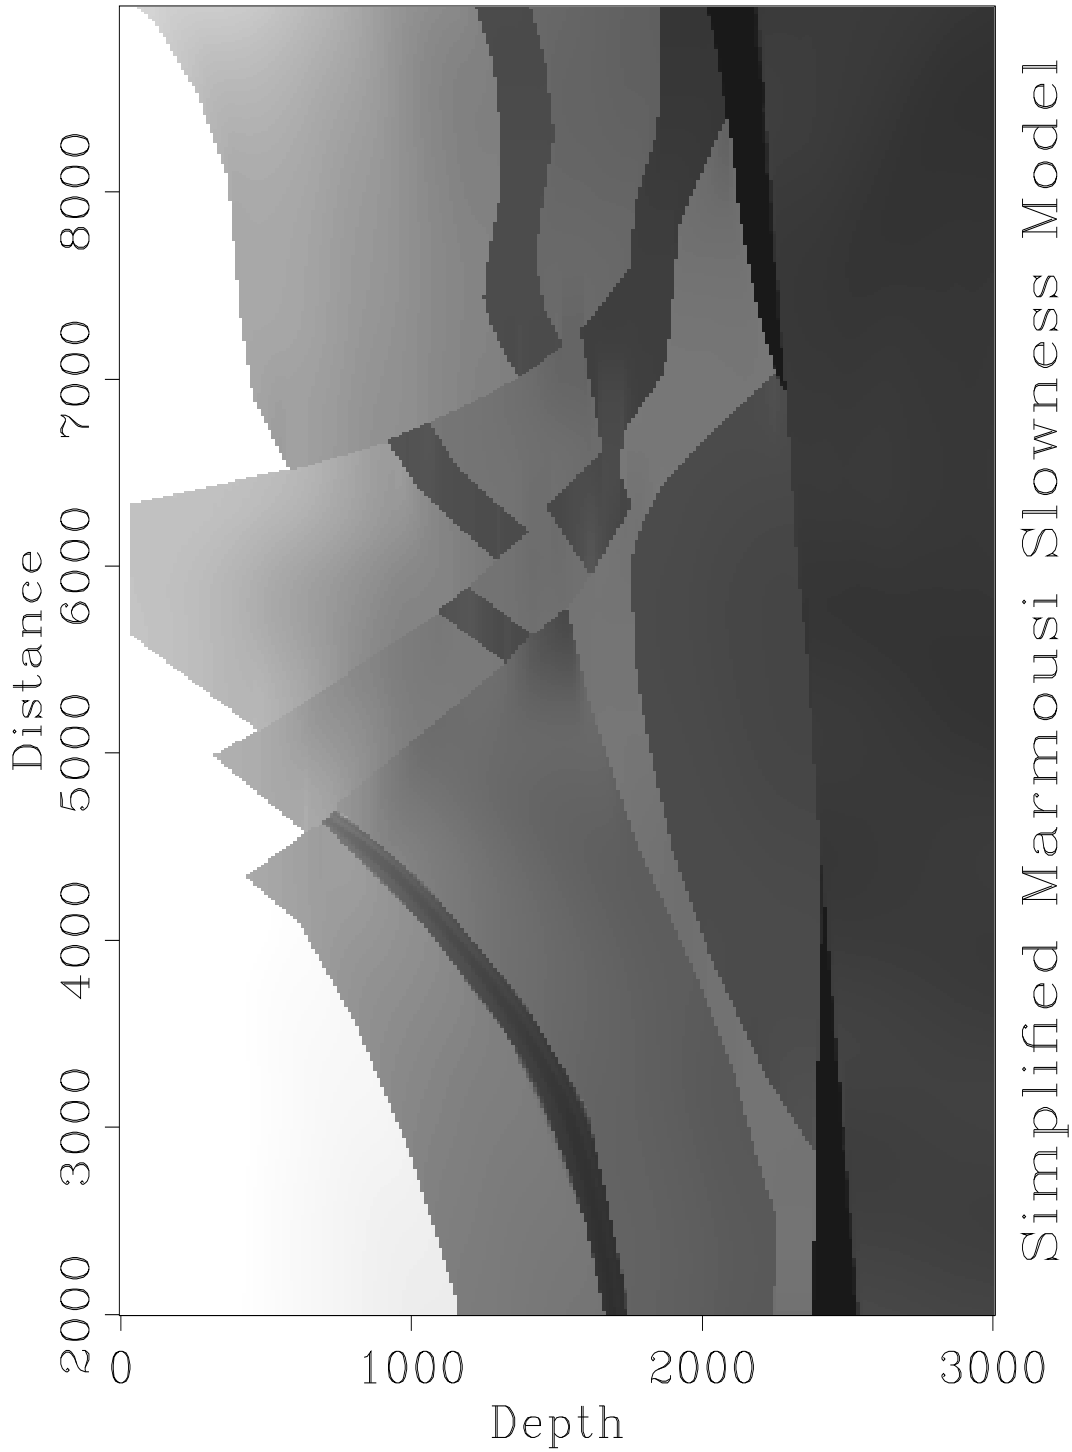

Figure 1.18: A simplified version of the true slowness field. The slowness was smoothed within structurally inspired regions. MigMarm-geol-slow [CR]

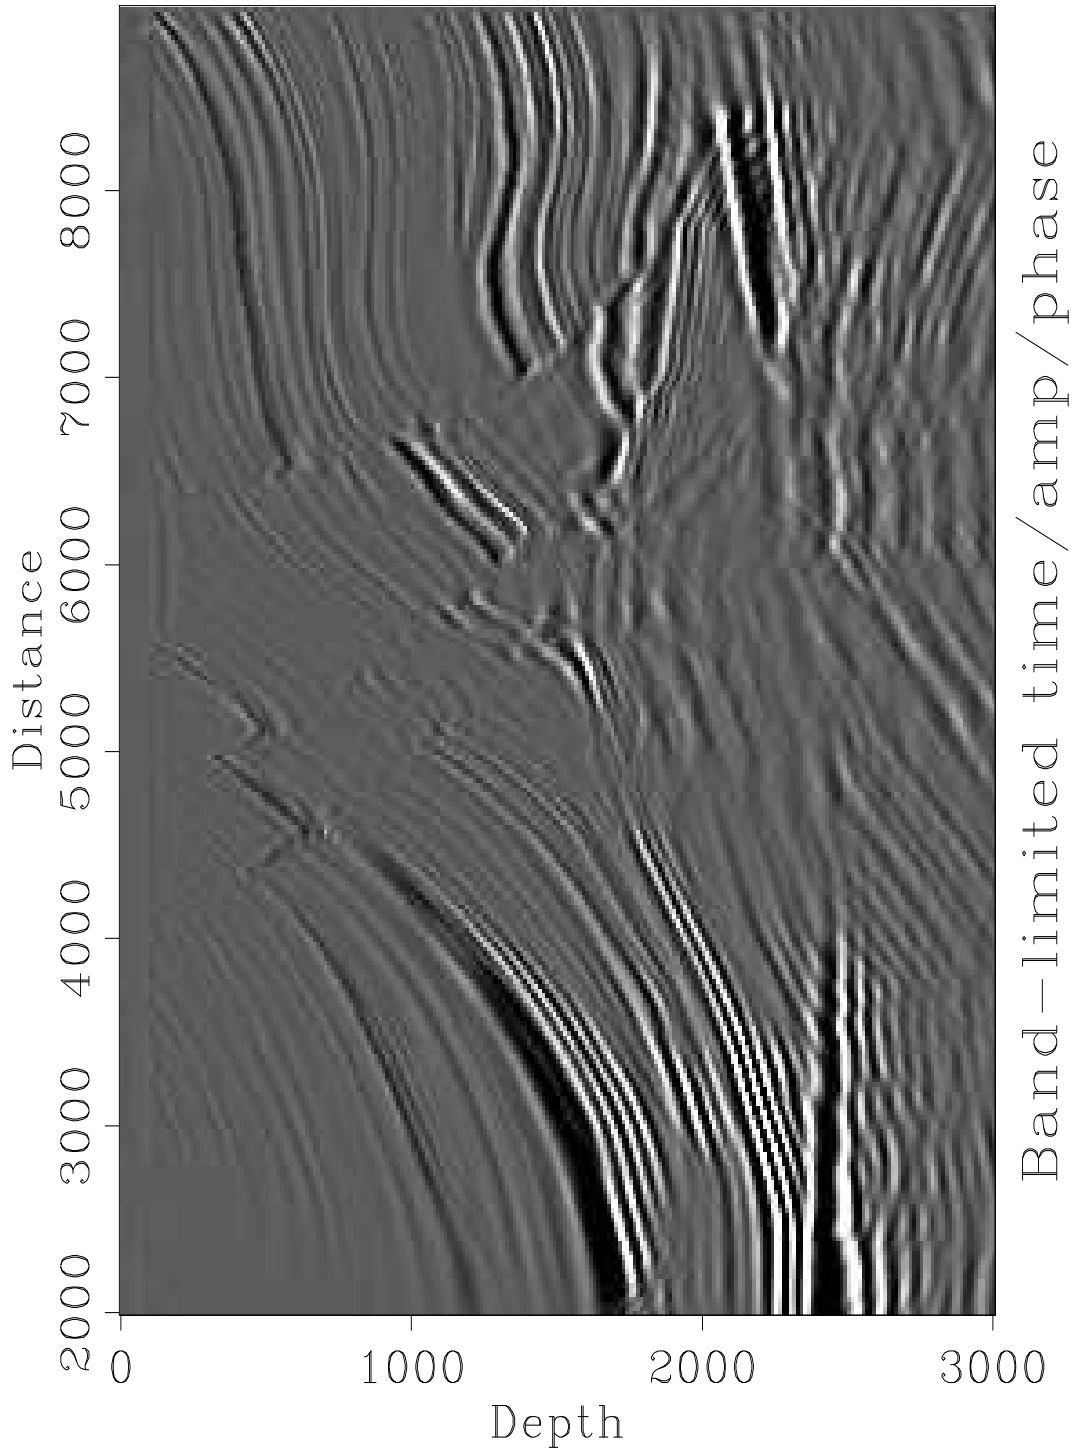

Figure 1.19: Marmousi dataset migrated using band-limited traveltime, amplitude, and phase maps calculated in the simplified model. Compare with Figure 1.13. MigMarm-geol-tap [CR]

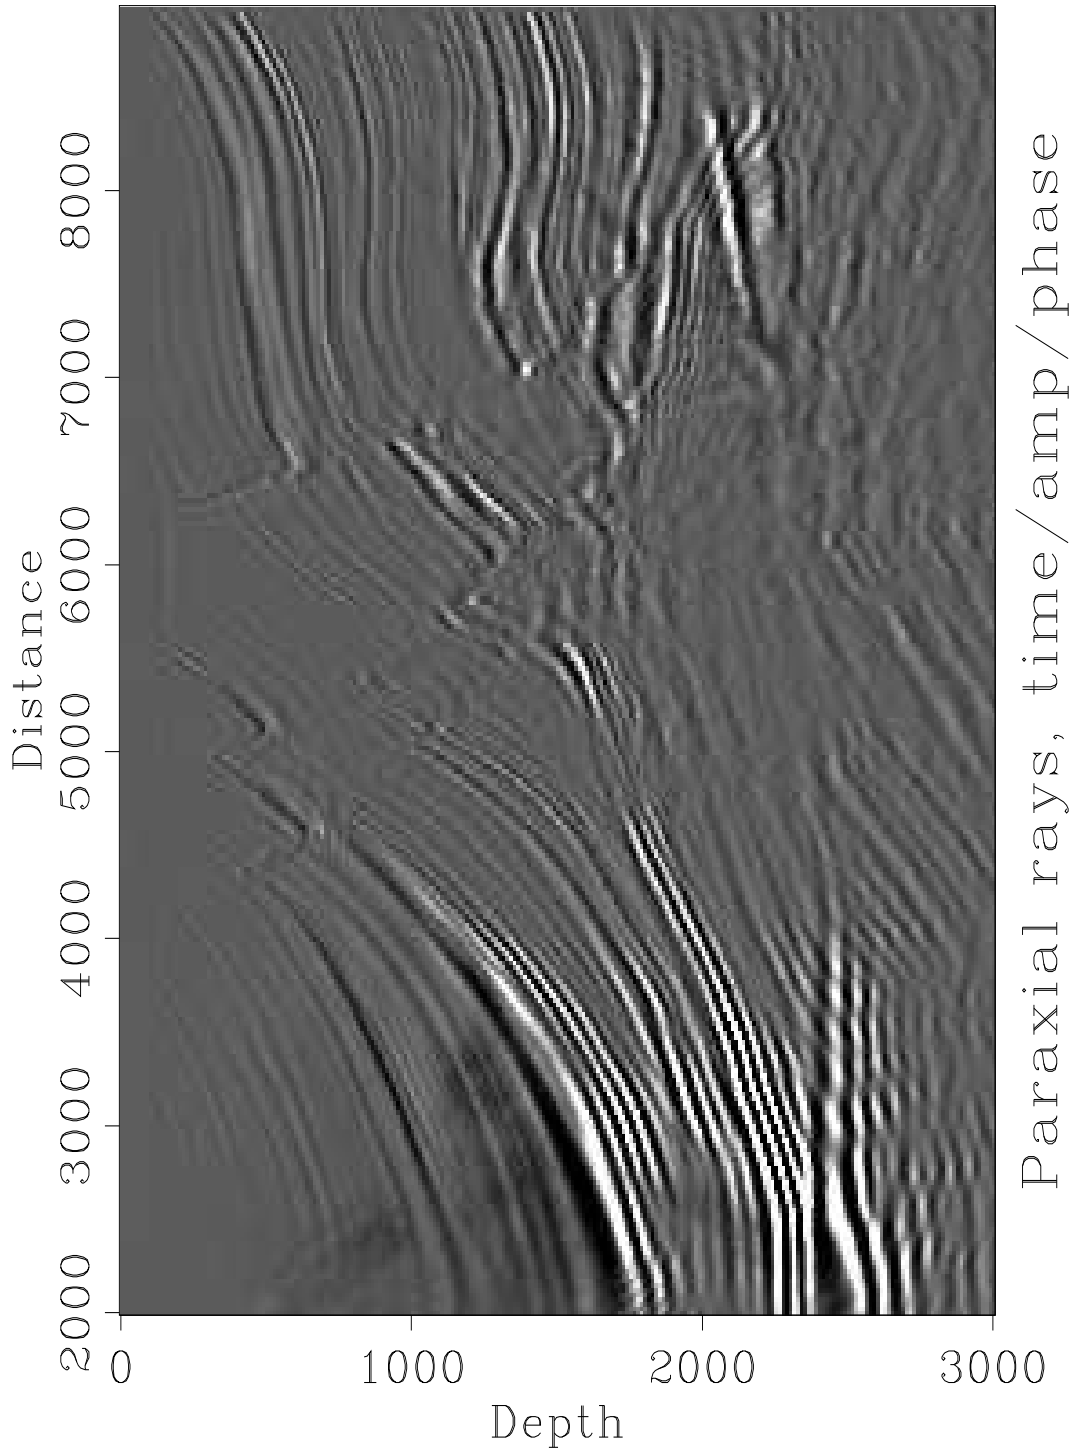

Figure 1.20: Marmousi dataset migrated using paraxial ray tracing traveltime, amplitude, and phase maps calculated in the simplified model. Compare with Figure 1.9. MigMarm-parax-simp-tap [NR]

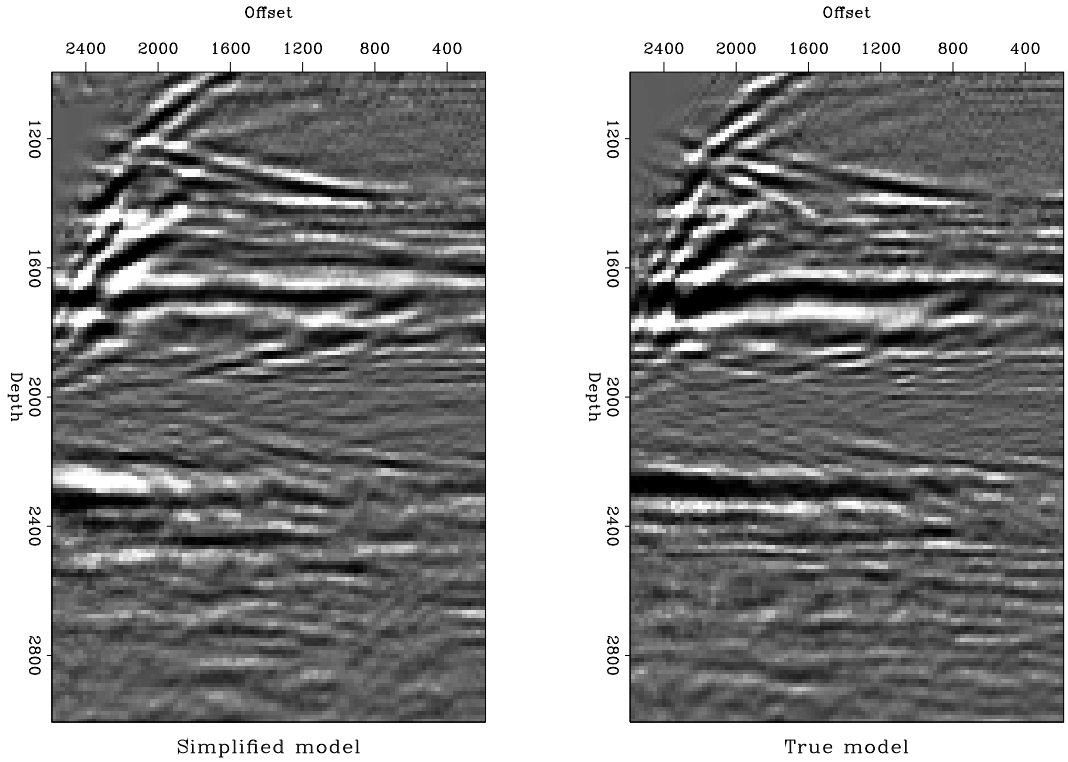

Figure 1.21: Comparison of CDP 7000 after imaging with the simplified model and the true model. MigMarm-comp-cdp [CR]

to give a poor image. Images created using the maximum-amplitude traveltime, calculated by paraxial raytracing, are much better than those created using first-arriving traveltimes.

The band-limited Green's functions calculated using my method give a structural image that is even better than the one created using paraxial raytracing traveltimes. My result is very close to the output of finite-difference shot-profile migration. My method has several advantages over ray-based methods. It is calculated on regular grid, not in ray coordinates, it does not use an asymptotic approximation, and it is unconditionally stable. No smoothing of the slowness model is necessary to stabilize the algorithm.

Figure 1.22 compares the four methods in the Marmousi model. All of these images are structural images, not an attempt to recover the true reflectivity. The results of finite-difference shot-profile migration and Kirchhoff migration using band-limited Green's functions are the two clearest images. They are very similar in quality but the Kirchhoff migration is cheaper to compute. Kirchhoff algorithms have other added advantages; they can be used when spatial sampling is irregular, and they can be used to image a target that is a subset of the whole volume without having to image everything above the target level. Both of these advantages become much more significant when migration of 3-D prestack data is attempted. On current computers full 3-D prestack finite-difference shot-profile migration is not feasible, but target-

oriented 3-D prestack Kirchhoff migration is definitely practical and is in wide use. I have shown a method for calculating Green's functions that should permit high quality images even when a complex velocity model is used.

If the correct imaging condition is used, we can use a Kirchhoff migration/inversion to recover a better estimate of the relative reflectivity of the subsurface. Figure 1.22 shows a comparison of the ideal result in the target zone, with the results obtained from paraxial-ray tracing and band-limited Green's functions. I believe that the band-limited Green's functions give a better image because they estimate the effects of wave propagation in the seismic bandwidth, not wave propagation at very high frequency.

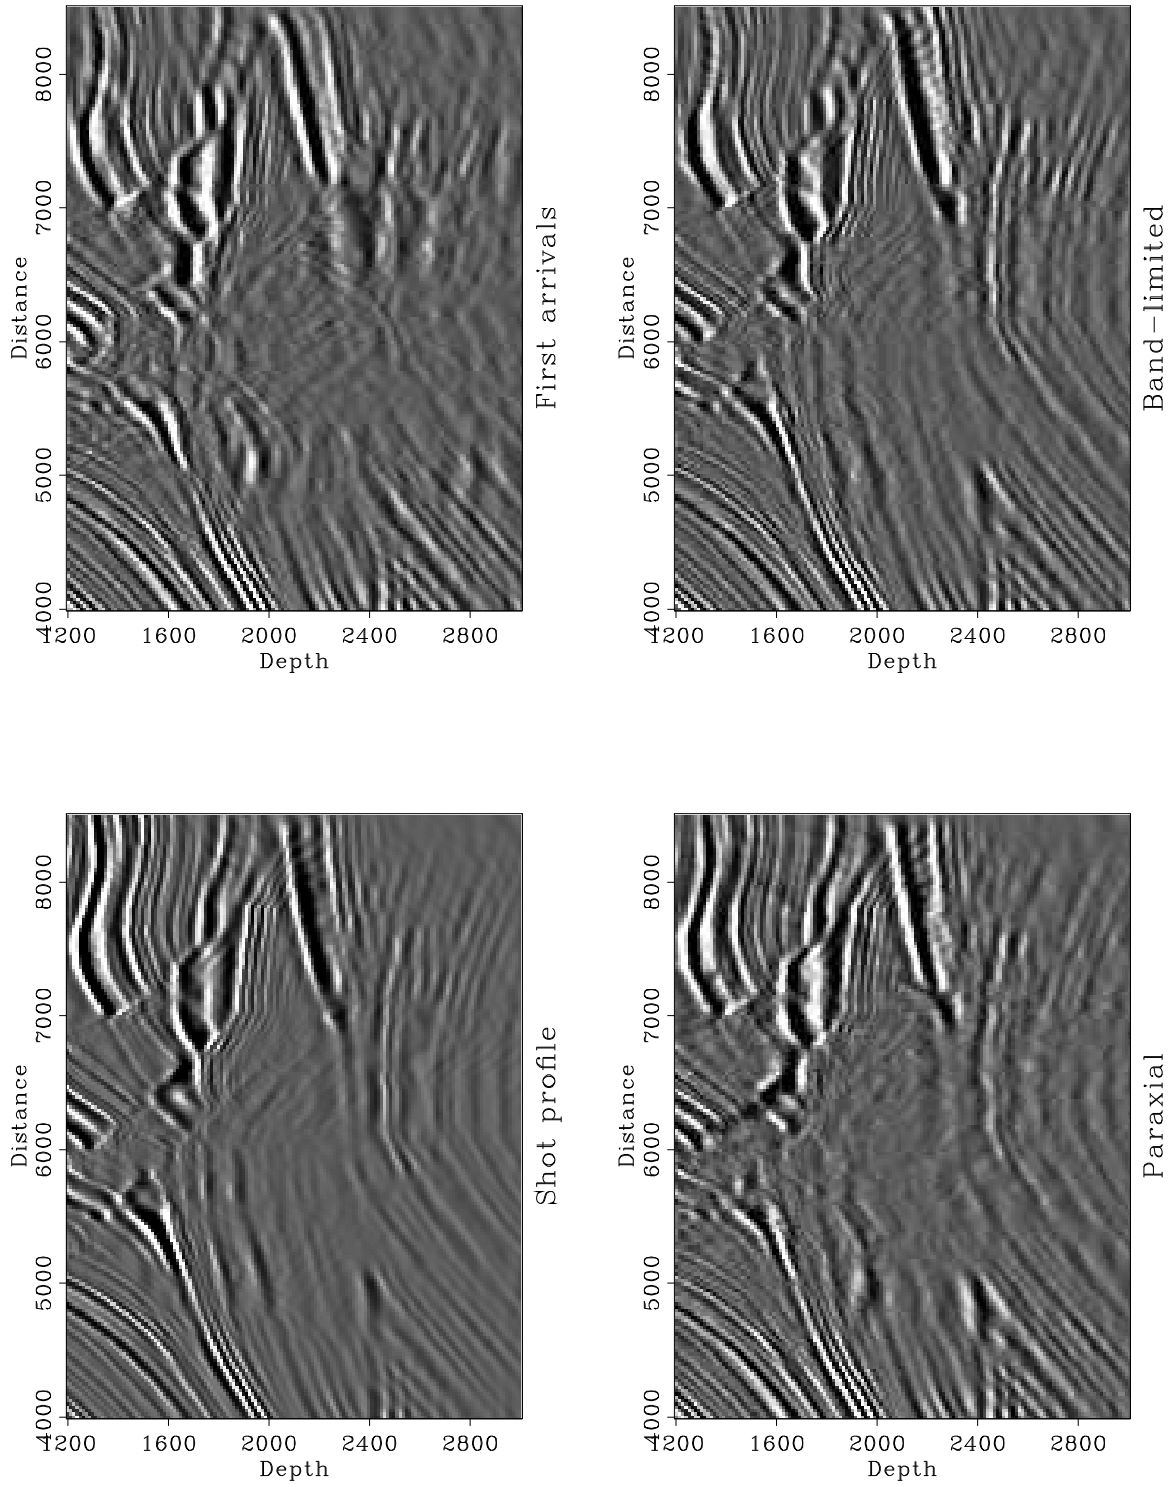

Figure 1.22: Comparison of structural images. MigMarm-migex2 [CR]

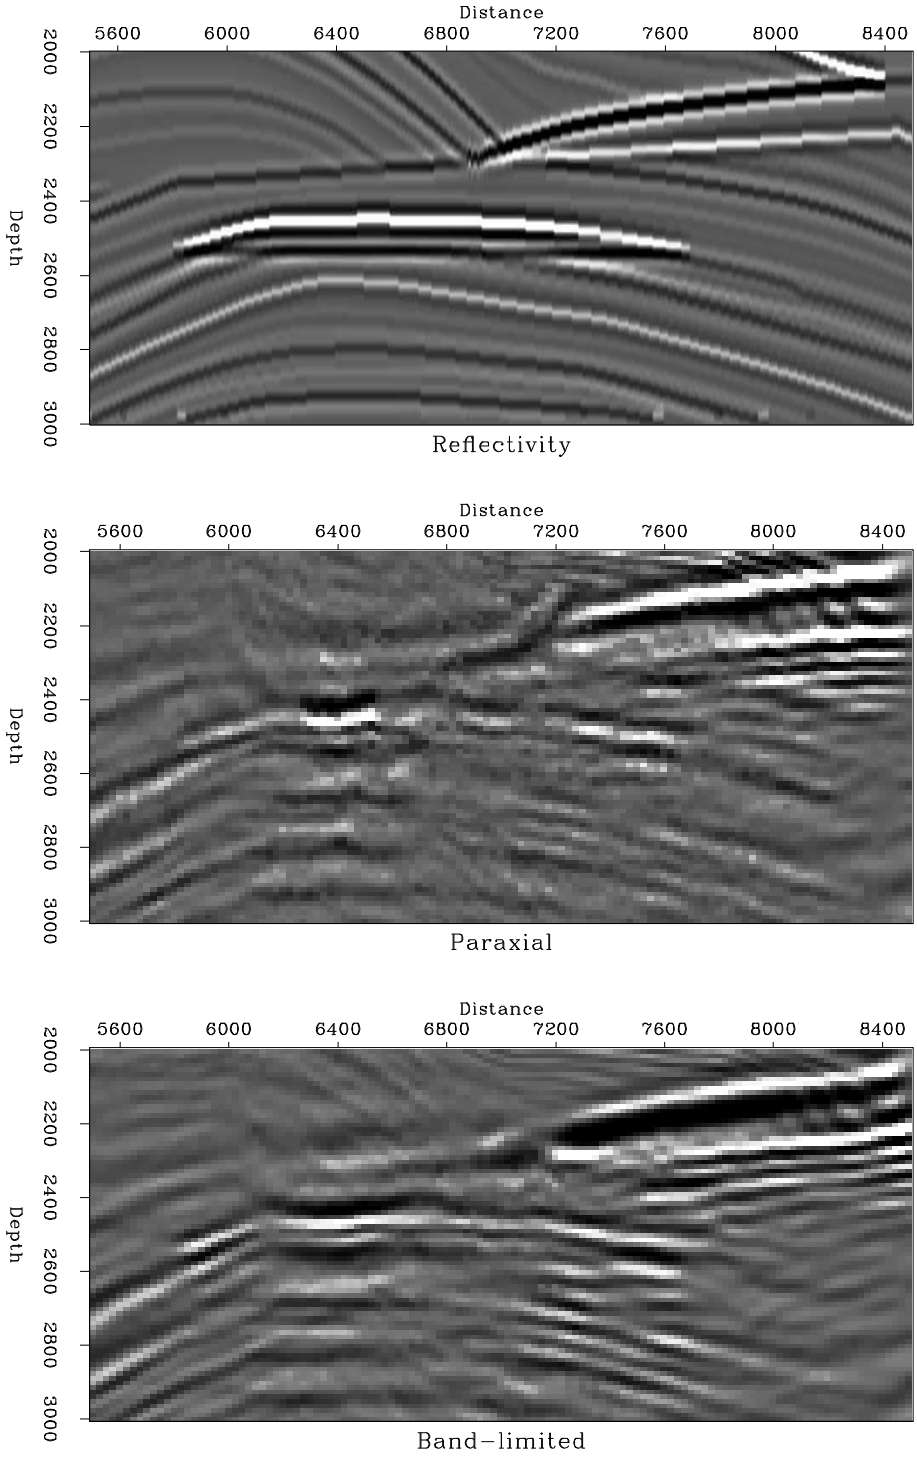

Figure 1.23: Comparison of reflectivity images. Top is the ideal image. Middle is the image from paraxial ray tracing. Bottom is the image created using band limited Green's functions. MigMarm-migex3 [CR]