Chapter 1

Band-limited Green's functions in the Marmousi model

In this chapter I study different ways of calculating Green's functions in the Marmousi model. I compare traveltimes and amplitudes computed using the band-limited traveltime method, with full one-way and two-way wavefields, and with traveltimes calculated by finite difference solutions of the eikonal equation and paraxial ray tracing. Figure 1.1 is a plot of the slowness field for the model. It has a complicated faulted structure that impedes imaging of the anticlinal trap at the bottom of the model.

1.1 Comparison of band-limited and first-arriving traveltimes

Figure 1.2 shows contour plots of three sets of traveltimes superimposed on the Marmousi velocity field. At the top are the first arriving traveltimes, calculated by a finite difference solution to the Eikonal equation. In the middle are maximum amplitude traveltimes from paraxial ray tracing. At the bottom are band-limited traveltimes, calculated using my method. The initial impression is that the first arrival traveltimes are much more reasonable (at least they are much prettier) than those calculated using my method. My task in this chapter will be to persuade you otherwise; pretty isn't always right! The results from paraxial ray tracing are closer to the results from my method than to the first arrival traveltimes. This is to be expected as they are both methods are estimating the maximum amplitude arrival rather than the first arrival. There are still significant differences between my method and paraxial ray tracing. In some regions they are choosing different wavefronts as the maximum amplitude event.

Figures 1.3 and 1.4 show time-slices through the full outgoing wavefield with traveltime contours superimposed. The contours of the band-limited traveltimes are a much better match to the significant amplitudes in the wavefield snapshot than the first arriving traveltimes. The paraxial traveltimes also appear to be a good

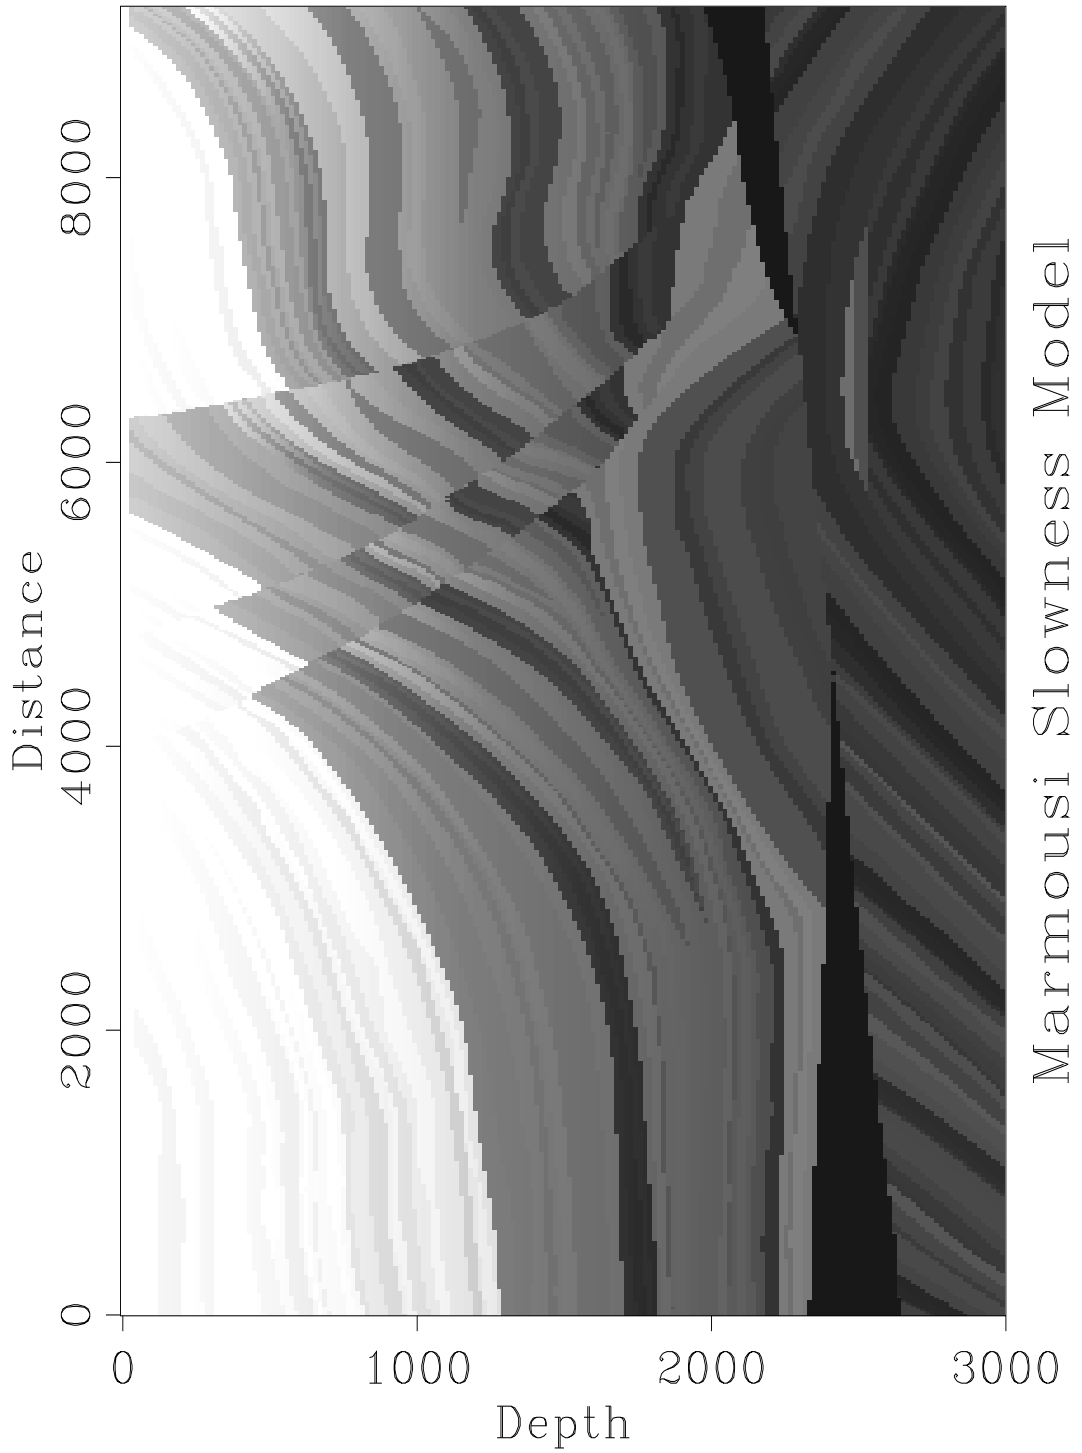

Figure 1.1: The slowness model for the Marmousi dataset. MarmTT-marmslow [ER]

match even though they choose different events in some parts of the model. Both the paraxial ray tracing and the band-limited Green's functions have a phase associated with each point, in addition to a traveltime. If these methods choose a phase shift of $\pi/2$, the traveltime contours should follow the zero crossing rather than the peak of the wavefront. In the lower parts of the model they are both choosing a non-zero phase event. This will be more clearly shown later in this chapter. The finite difference solution to the eikonal equation always estimates a zero phase arrival so the traveltime contours should always follow the peak of the leading wavefront.

The full wavefield has many events present in it, but the band-limited traveltime scheme (as currently implemented) only chooses one event at each location. This accounts for the large jumps in the band-limited traveltime contours. The algorithm is switching between events in order to choose the maximum energy event at each location. While this is a limitation, it does have the property that the event chosen is always the most significant one. In contrast, the first arriving traveltime field is always continuous but there is no guarantee that the events contain any significant energy. In particular the first arriving wavefront in the left-center of the figures is clearly one with very little energy. The reason for this can be clearly seen in figure 1.2. That portion of the wavefield is composed of head waves from the thin high velocity layer that runs from (x, z) coordinates (2000,1700) to (4000,1200). The only energy in this part of the wavefront comes from the wave that enters the top of the thin layer.

The least convincing area of the band-limited traveltime contours is in the bottom left of the plots. The contours have some wild swings in them. At some locations these contours even appear to be slightly ahead of the first arriving traveltime. This area is one which almost no energy reaches from the given source location. My algorithm is less reliable when all the events in the picking window are very weak. This should not be a problem if these traveltimes and amplitudes are used for a Kirchhoff migration algorithm, since if the large amplitudes are correctly accounted for then the mispositioning of a few very weak amplitudes will not be significant.

In order to have a qualitative understanding of the effect that these traveltime differences would have on a Kirchhoff migration algorithm, we can look at slices of the traveltime fields and the full wavefield. In the following figures a wavefield was modeled from a point in the reservoir to the surface. The traveltimes from the surface to that point in the reservoir are superimposed on the wavefield. In a constant velocity field the traveltime curve would be a hyperbola and the wavefield would consist of a single band-limited hyperbolic event. Figures 1.5 to 1.7 show a very different picture for the Marmousi model. The wavefield has many non-hyperbolic arrivals, and the different arrivals cross, overlap, and merge. If we wish to image this data at the surface the Kirchhoff migration must integrate along a trajectory that passes through the significant energy in this wavefield.

The trajectory in Figure 1.5 is that given by a finite-difference solution to the eikonal equation. It has successfully defined the trajectory of the first arriving energy. Unfortunately, the first arrivals are very weak and this is a poor choice for imaging.

Paraxial ray tracing produces the curve shown in Figure 1.6. It does a much better job of following the significant energy. However, is some high frequency oscillation in the traveltimes which may degrade the coherent summation of energy in the migration program. In contrast, the band-limited traveltimes in Figure 1.7 are smoother in the regions of high energy. In regions where little energy is present the picks can oscillate wildly, but this should not be important in imaging as long as the significant energy is accounted for correctly.

It should be noted that the last two figures display just the traveltime parameter for Green's functions that also have an associated phase parameter. This means that they may not always follow the peaks of the wavefield. If a 90° phase shift is associated with one location we would expect the traveltimes to follow a zero crossing in the wavefield. The best test of the combined traveltime and phase fields is to use them in a migration program. In the next chapter I show the results of Kirchhoff migration on the Marmousi synthetic dataset using the different estimates of the Green's functions.

1.2 Comparison of the full wavefield with the maximum energy wavefield

The band-limited traveltime/amplitude/phase triplets define a parametric approximation of the full wavefield. In fact two approximations have been made.

- The full acoustic impulse response has been approximated by the one-way, outgoing acoustic impulse response calculated in polar coordinates.
- The outgoing wavefield has been approximated by a simple, single event model at each location that is parameterized by a traveltime/amplitude/phase triplet.

Figures 1.8-1.10 each display time-slices of three wavefields. The top frame is the full acoustic impulse response calculated using a explicit,time-marching, finite difference scheme. This scheme correctly models the full wavefield. It includes primary and reflected arrivals, multiple reflections, etc. The middle frame is the wavefield calculated by one-way extrapolation in polar coordinates. All of the frequencies were extrapolated and a time domain response was calculated by inverse FFT. The bottom frame is the wavefield defined by the traveltime/amplitude/phase triplets at each location. The triplets were used to calculate a frequency domain representation, and this was the inverse Fourier transformed to give the time domain wavefield.

If the parametric representation of the impulse response is a good approximation to the full wavefield then we would expect that migration with the parametric Green's function will give satisfactory results. If the parametric representation does not correspond to the major events in the full wavefield, then we might expect the migration to fail.

The two-way wave equation modeling result has many more events within the outer wavefront than the outgoing wave equation modeling result. The two-way

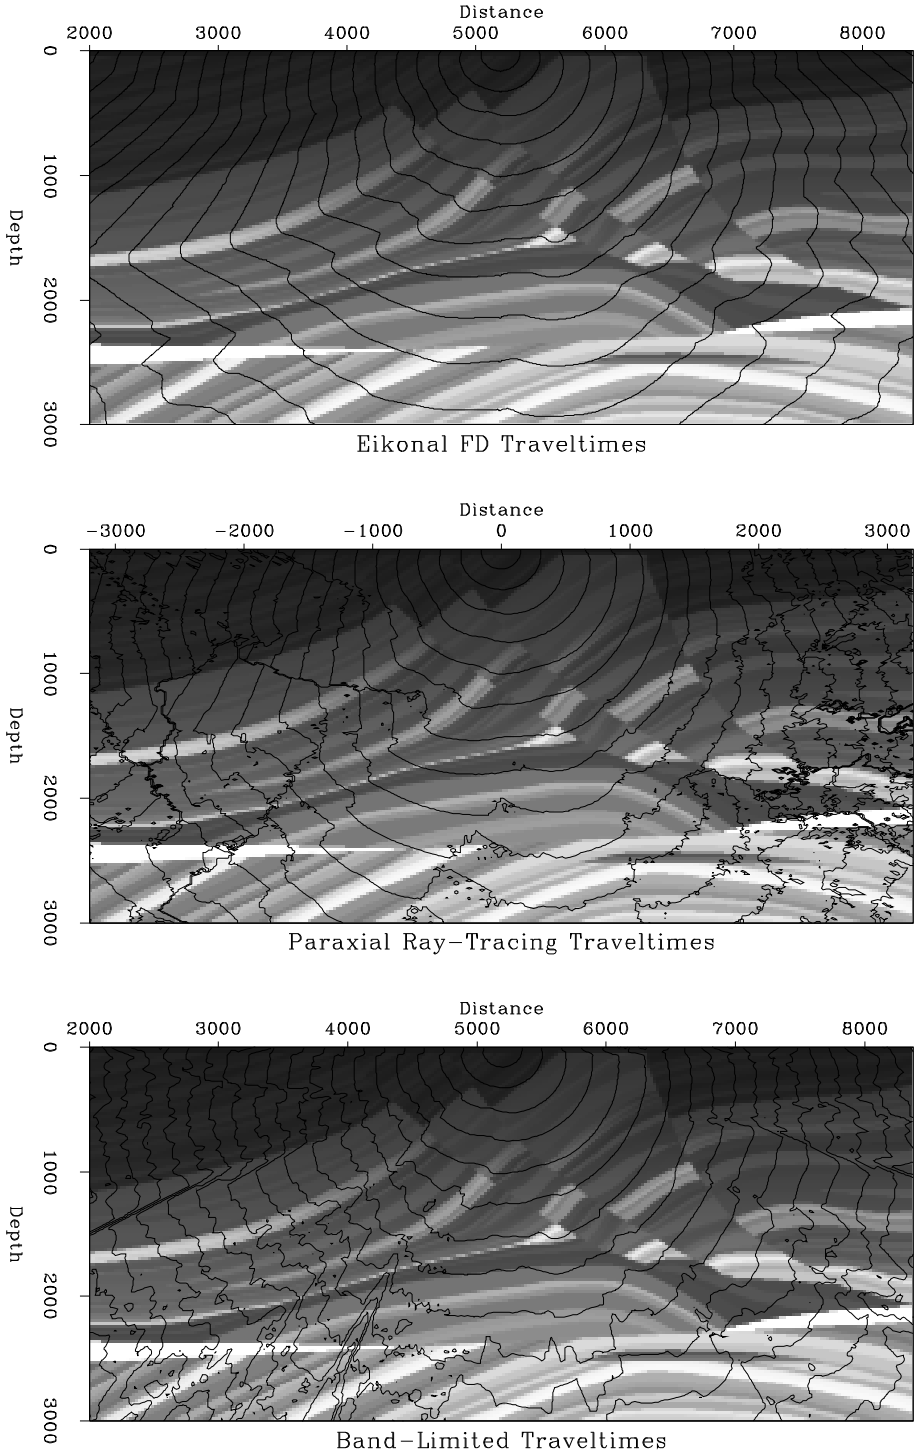

Figure 1.2: The top figure shows first arriving traveltime contours calculated from the eikonal equation. The middle figure shows maximum amplitude traveltimes calculated by paraxial ray tracing. The bottom figure shows band-limited traveltimes calculated from the wave. The contours are plotted at 0.1 sec. intervals MarmTT-marm2-tt [ER]

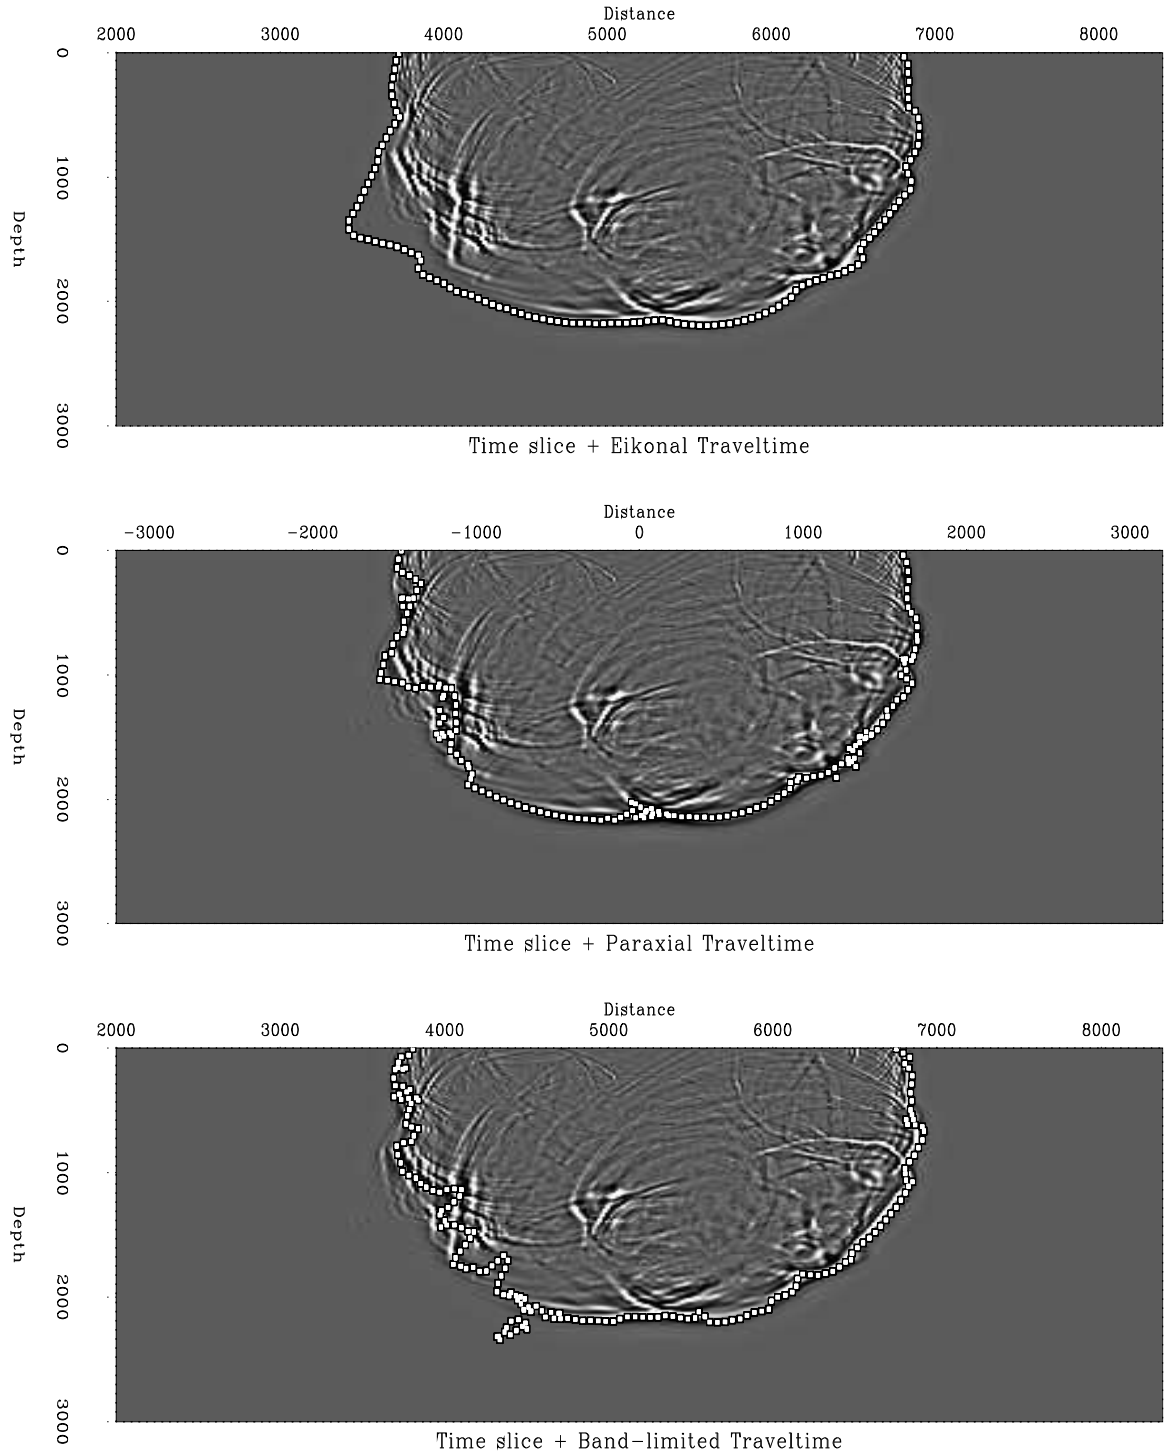

Figure 1.3: A snapshot of the full wavefield at 0.9 sec. with the 0.9 sec. contour of the traveltime field superimposed. Top, finite difference solution to eikonal. Middle, paraxial ray tracing. Bottom, band-limited traveltimes. MarmTT-all-tov-.9 [ER]

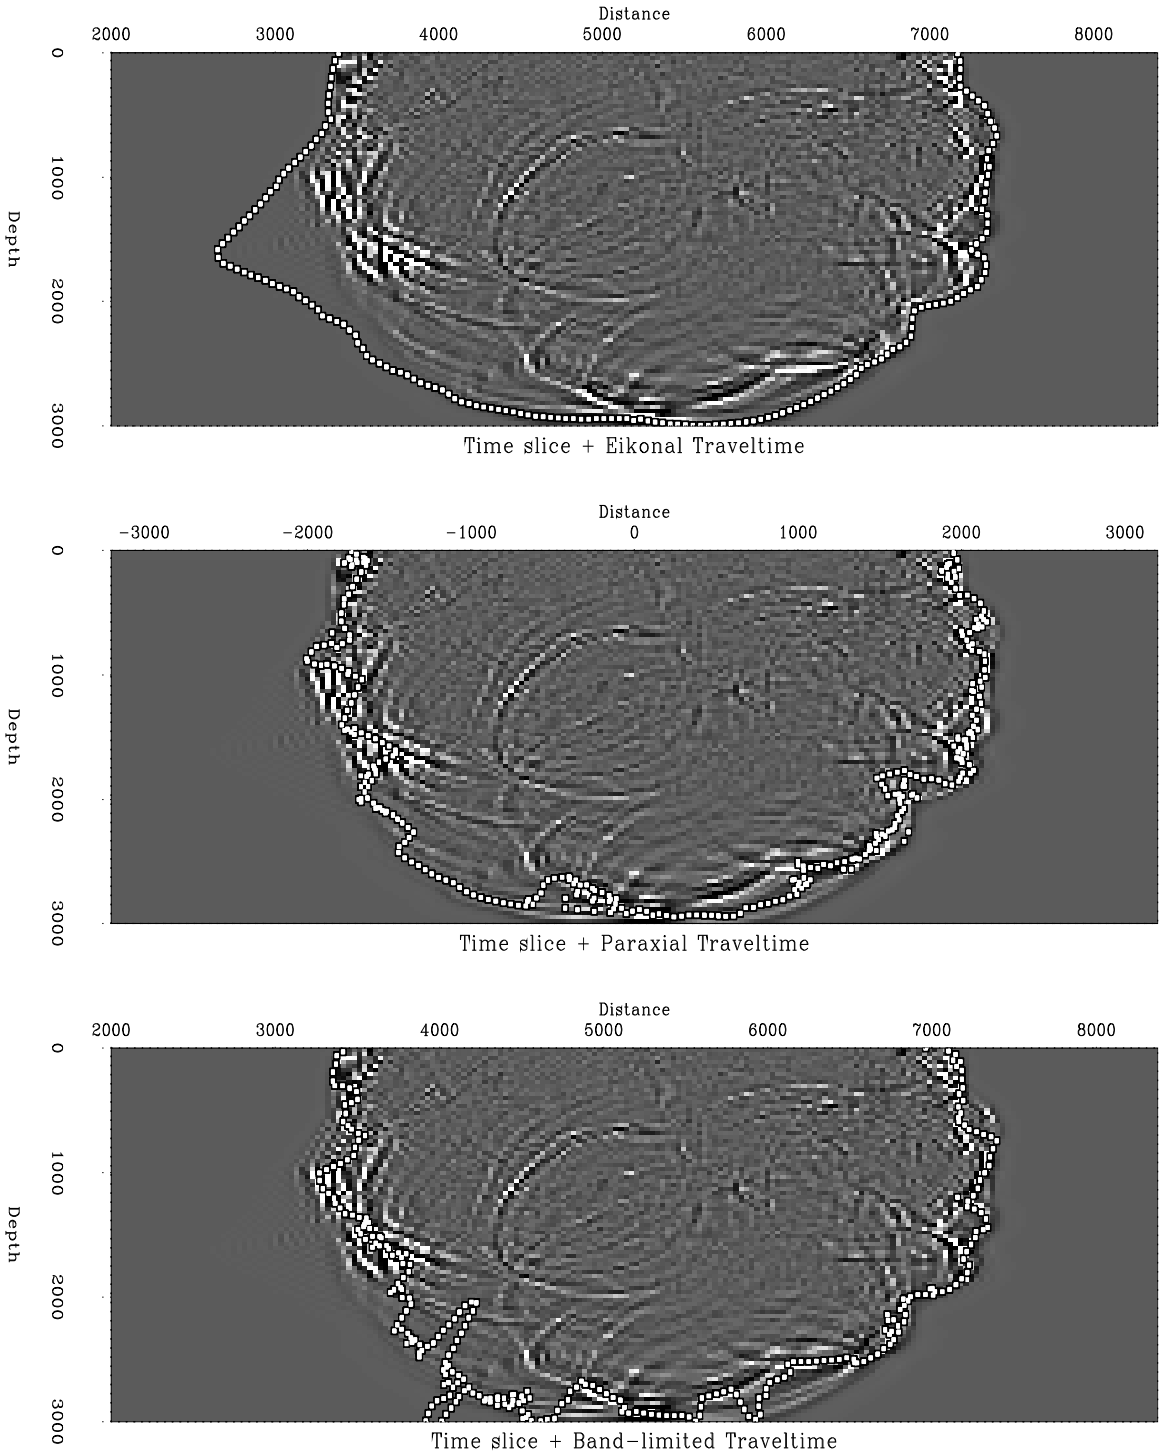

Figure 1.4: A snapshot of the full wavefield at 1.1 sec. with the 1.1 sec. contour of the traveltime field superimposed. Top, finite difference solution to eikonal. Middle, paraxial ray tracing. Bottom, band-limited traveltimes. MarmTT-all-tov-1.1 [ER]

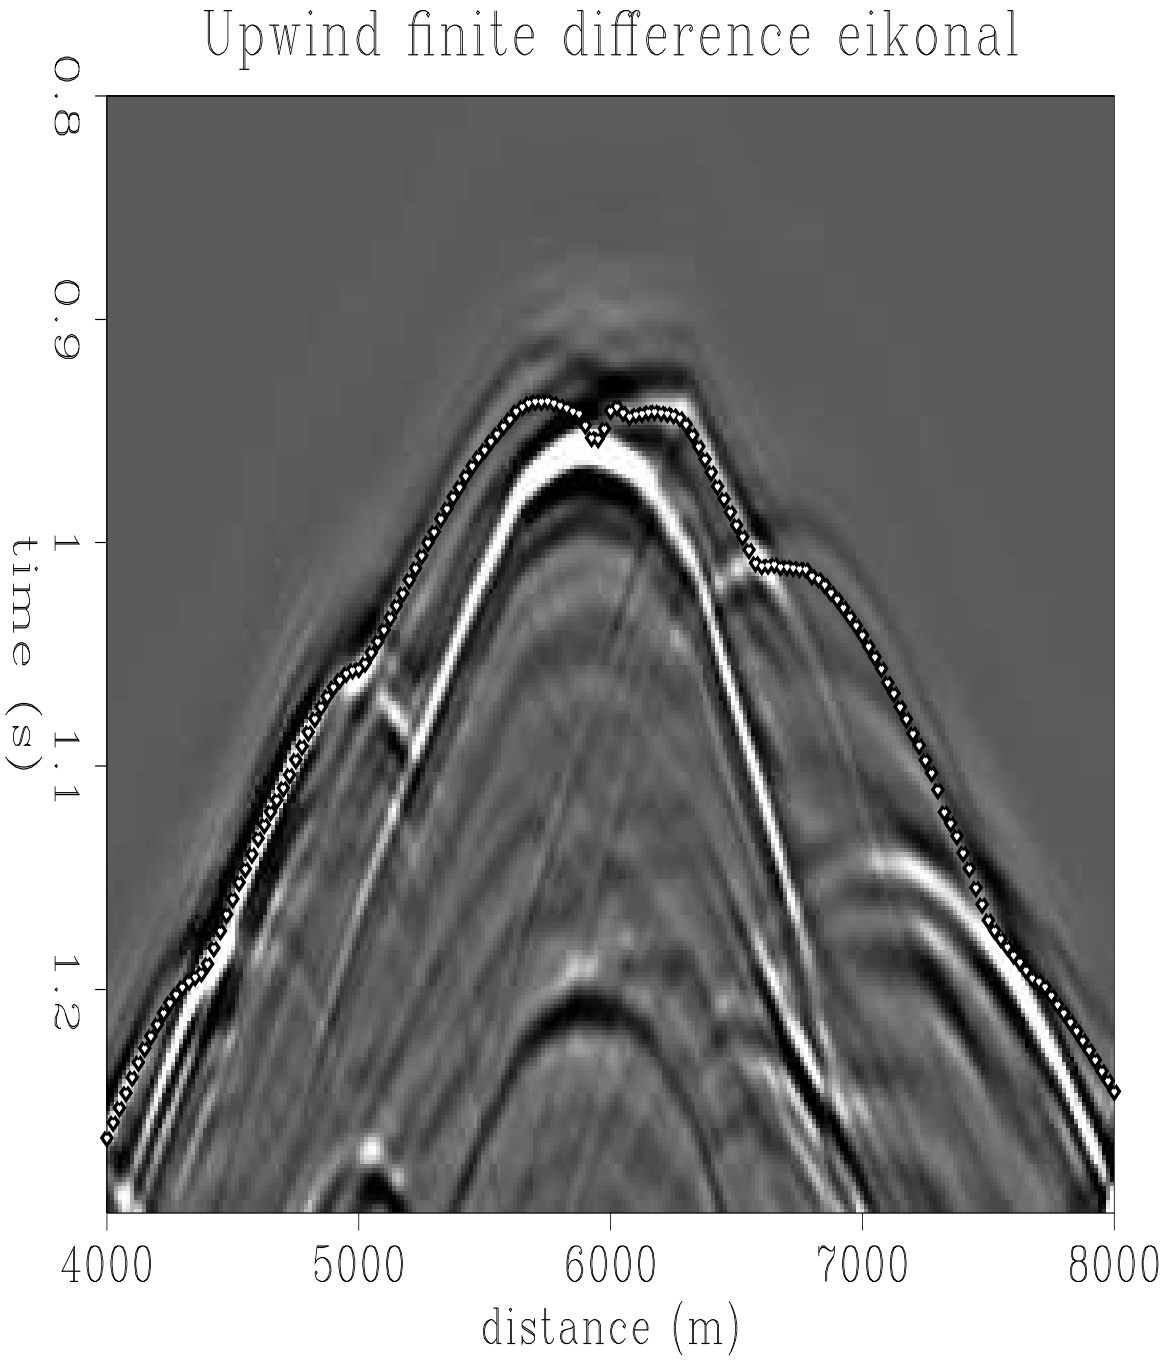

Figure 1.5: The full wavefield from a source in depth with first arriving traveltime picks superimposed. MarmTT-vantrier-tt [NR]

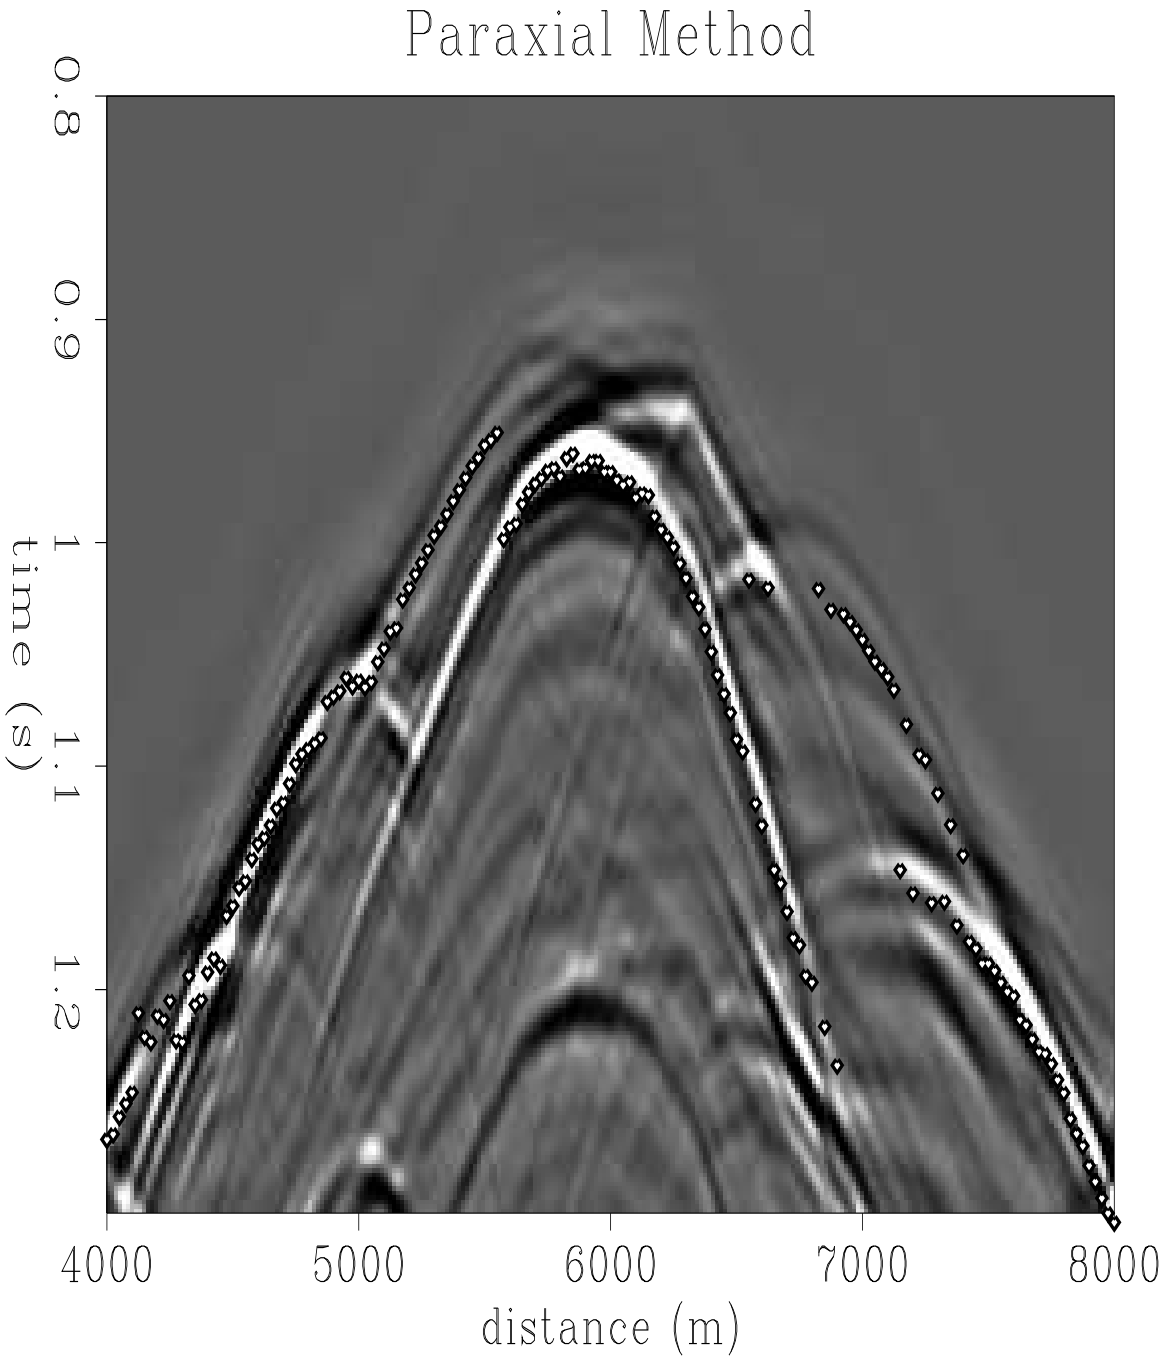

Figure 1.6: The full wavefield from a source in depth with paraxial ray tracing traveltime picks superimposed. MarmTT-paraxial-tt [NR]

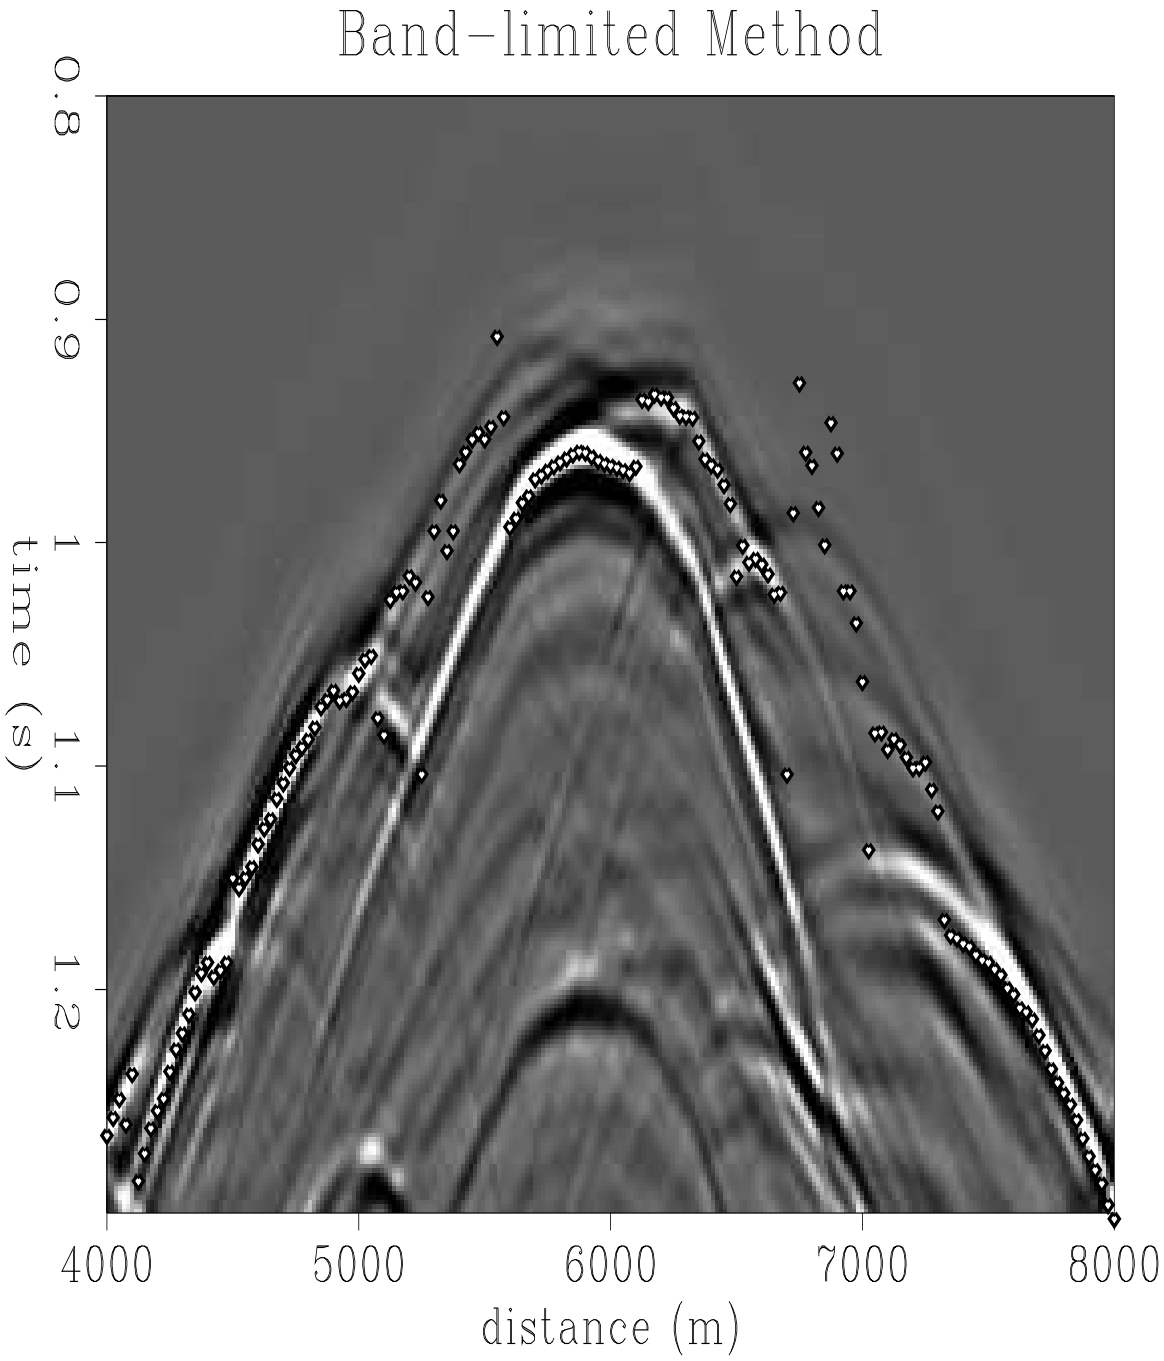

Figure 1.7: The full wavefield from a source in depth with band-limited traveltime picks superimposed. MarmTT-bandlim-tt [NR]

wave equation models multiples, and other back-scattered energy. The outgoing wave equation only models the forward scattered energy. These three figures show that the band-limited traveltime method produces a wavefront that is a good match to the first major arrival in the full wavefield. The traveltimes, amplitudes, and phases of all three wavefields match well. There are slight differences in timing and amplitude behavior between the two-way wavefield and the outgoing wavefield. These mat be due to the different initial conditions and boundary conditions in the two modeling schemes.

In areas of strong triplications, the band-limited impulse response only contains one branch of the triplication. This may degrade the imaging properties of this approximate Green's function. The scheme always picks the most energetic branch of the triplication. If only one event is picked at each location, the most energetic event is a good choice. The single event model is not the only parametric model that could be used to approximate the impulse response. I have chosen to restrict myself to a single event model in order to simplify the implementation of this algorithm. A multiple event model would require more complicated code but would not be intrinsically harder. A multiple event model would be a better match to the outgoing wavefield and thus a better match to the full acoustic impulse response. The utility of a multiple event model is probably very case specific. In the next chapter I will show that I can obtain satisfactory images of the Marmousi dataset using a single event model. I therefore chose not to implement a more complicated multi-event algorithm.

1.3 Comparison of paraxial ray tracing with band-limited traveltimes

A finite-difference solution to the eikonal equation only computes a traveltime field. In contrast both paraxial ray tracing and my method calculate traveltime amplitude and phase. The two particular implementations that I use both attempt to estimate the maximum energy arrival. Figure 1.2 showed that the two methods produce traveltimes that are similar but not identical. In Figures 1.3 and 1.4 they both appeared to be reasonable approximations to energetic parts of the full wavefield.

Figures 1.11 and 1.12 show amplitude and phase maps for the two methods. Although they have some gross features in common they are significantly different. The amplitude map from the band-limited Green's functions is much smoother than the map from paraxial ray tracing. In particular the wavefield passes through a caustic in the center of the plot. The asymptotic approximation used in paraxial ray tracing gives an infinite amplitude at the caustic. In this implementation the amplitudes have been thresholded to prevent this, however, there is still a clear spike in amplitude at the focus. In contrast a finite-bandwidth wavefront does not focus to a point and a much smoother, lower amplitude response is obtained. The largest region of difference between the two amplitude maps is the high amplitude region on the left of the paraxial ray tracing plot. I hypothesize that all of this high energy is due to one or two rays that have passed close to a caustic at coordinates (4900,1200). The

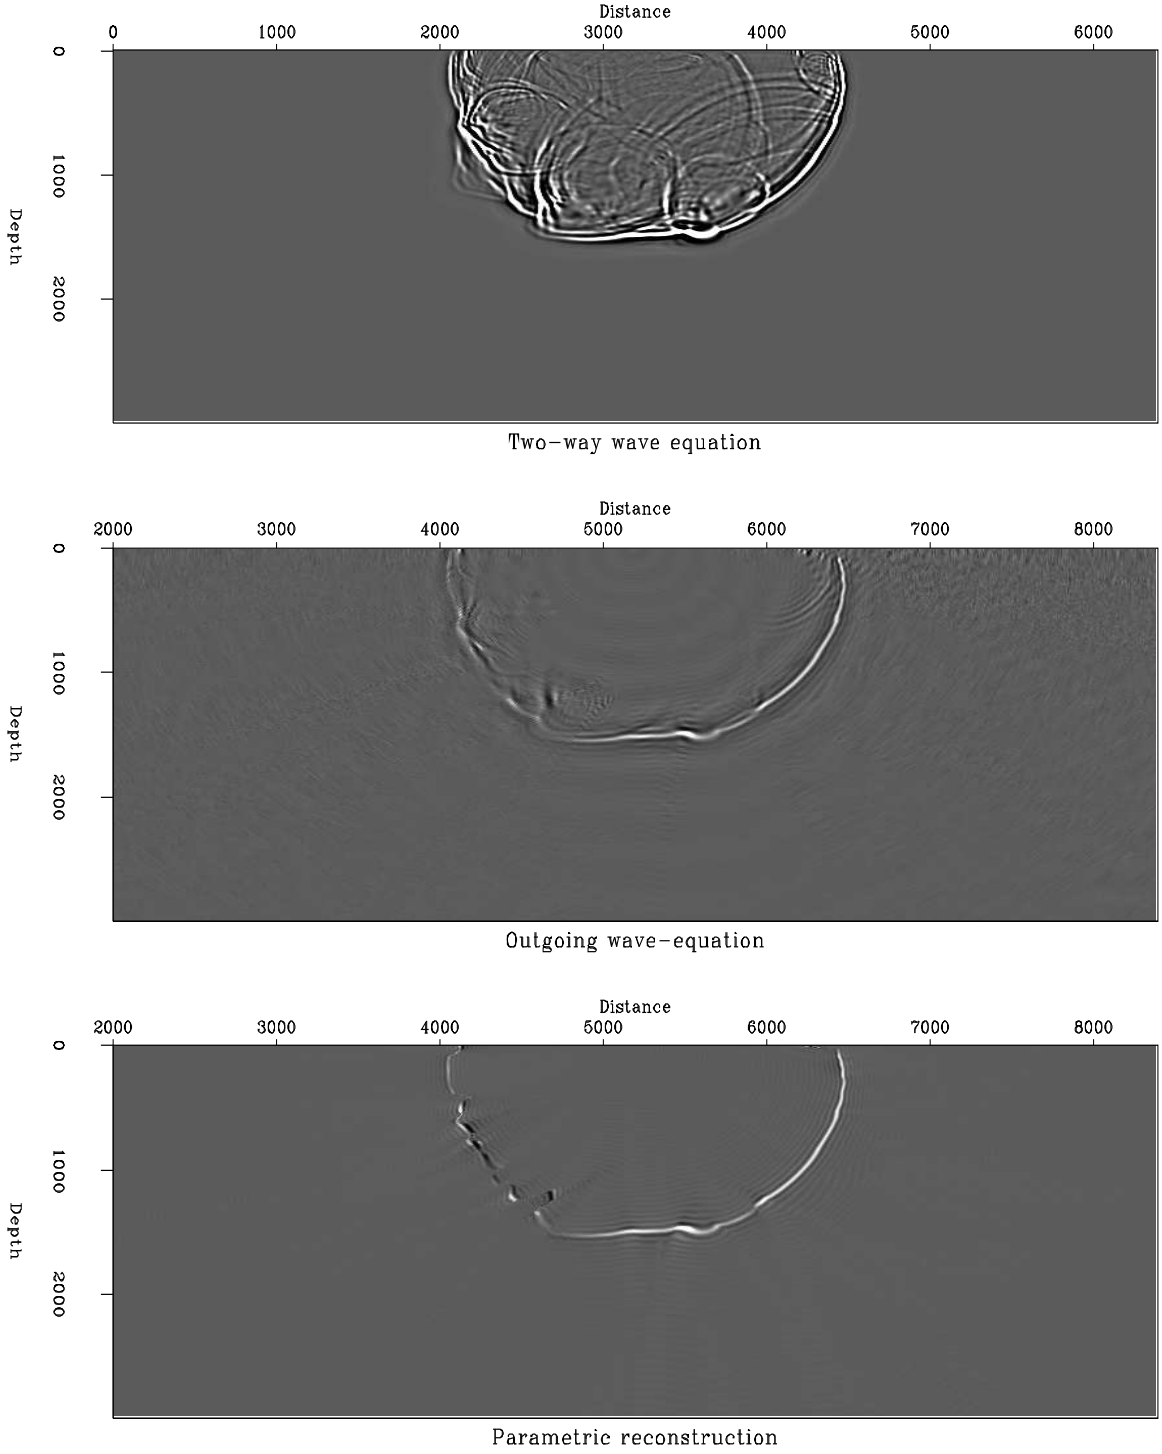

Figure 1.8: A time slice through the Green's function at 0.7 sec. The upper plot is the full two-way wavefield. The center plot is the full outgoing wavefield. The lower plot is the wavefield reconstructed from the band-limited traveltime/amplitude/phase fields. MarmTT-green-comp-.7 [CR]

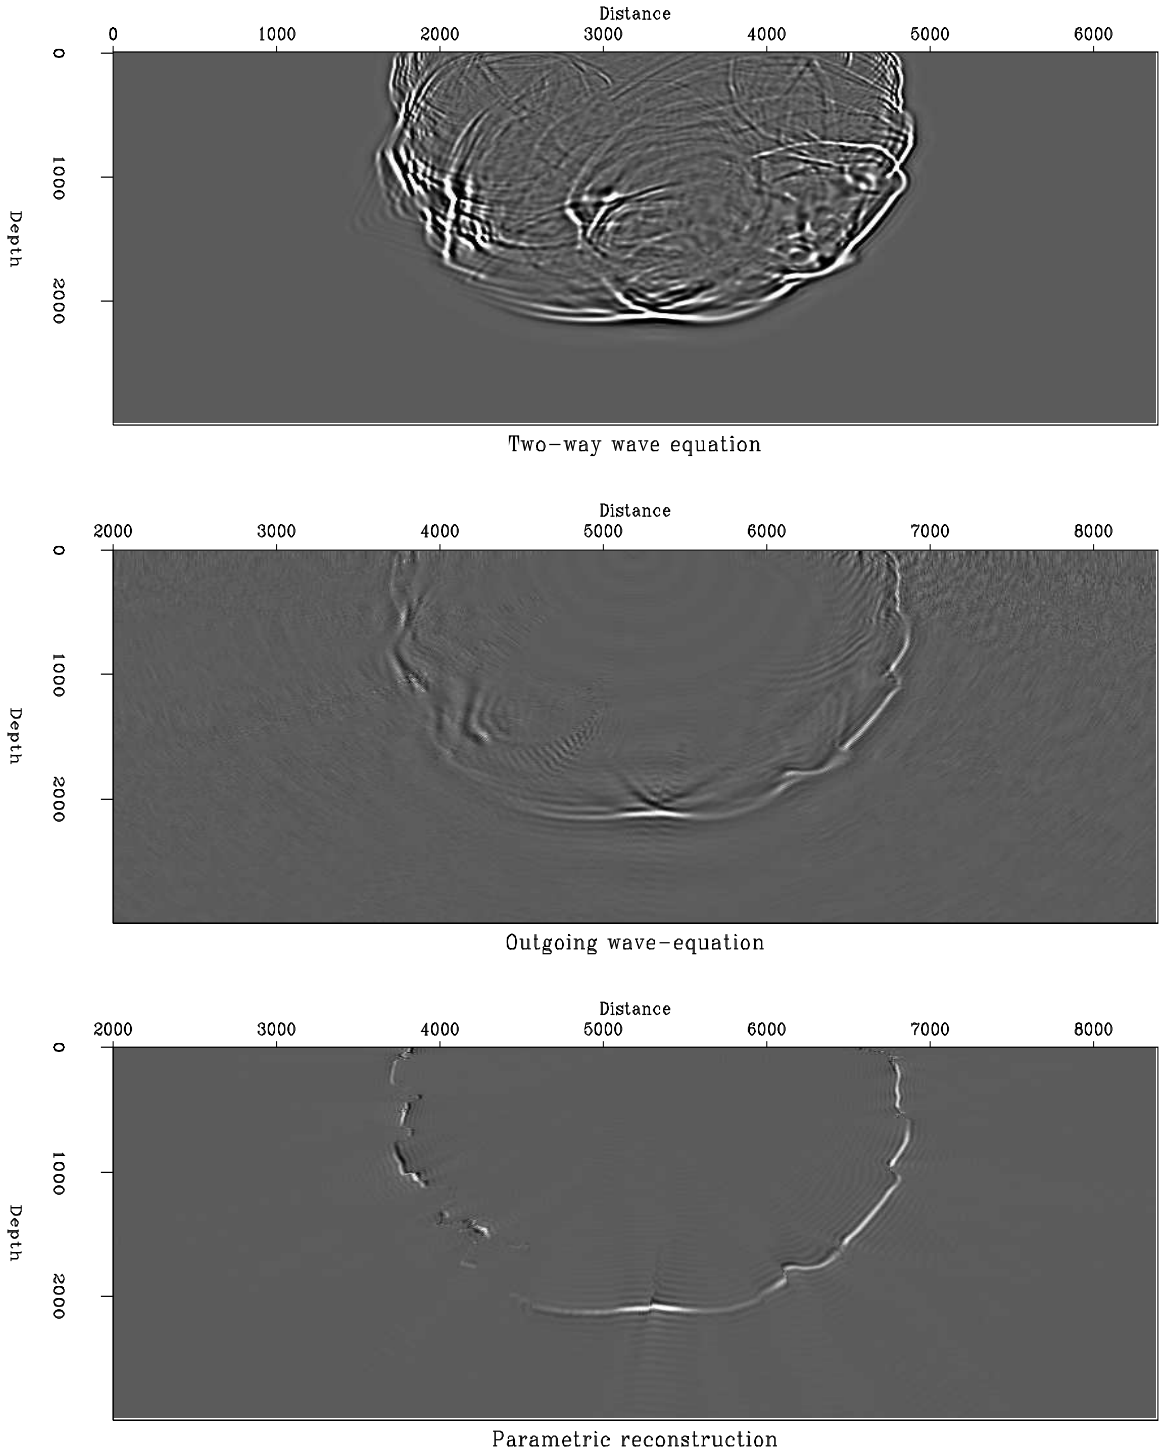

Figure 1.9: A time slice through the Green's function at 0.9 sec. The upper plot is the full two-way wavefield. The center plot is the full outgoing wavefield. The lower plot is the wavefield reconstructed from the band-limited traveltime/amplitude/phase fields. MarmTT-green-comp-.9 [CR]

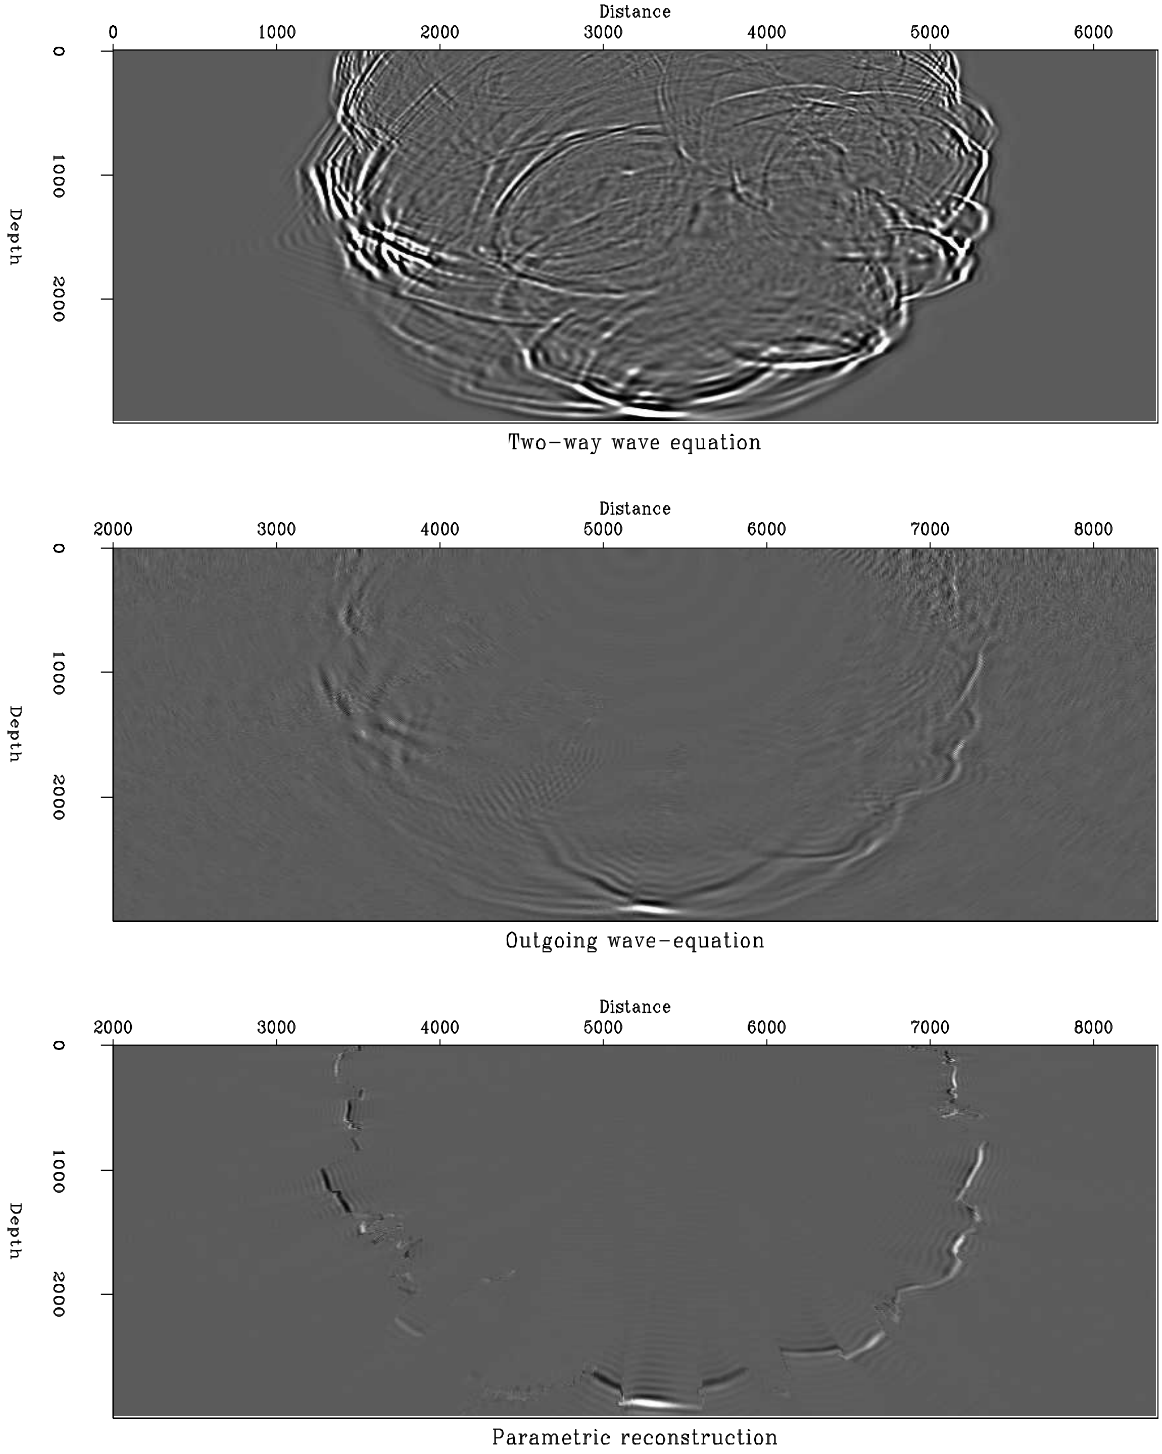

Figure 1.10: A time slice through the Green's function at 1.1 sec. The upper plot is the full two-way wavefield. The center plot is the full outgoing wavefield. The lower plot is the wavefield reconstructed from the band-limited traveltime/amplitude/phase fields. MarmTT-green-comp-1.1 [CR]

rays will then have an abnormally high amplitude. After the caustic the wavefront spreads out and this high amplitude is chosen as the maximum amplitude estimate over a wide region perpendicular to the raypath. These rays have become trapped between two high velocity layers and follow a zig-zag path to the edge of the model.

The phase plots shown in Figure 1.12 are also very different. The major reason for this is that paraxial ray tracing only calculates phases that are a multiple of $\pi/2$, giving the lower plot only four distinct shades. Each time a ray passes through a focus its phase is changed by $\pi/2$. Each ray in the ray tracing is associated with a single event. There is no mixing of phases between rays, so when a single ray is chosen to represent the maximum amplitude event it will always have a phase that is a multiple of $\pi/2$. The band-limited Green's functions may have any phase, as it is calculated directly from the modeled wavefield. This wavefield may be a superposition of two or more events. The phase is calculated from the complex wavefield at the peak of the energy function. Since the single event that is picked may be an approximation to multiple events, it is not unusual that the phase varies over the complete range. The rapid oscillation of the phase plot at large ranges is caused by the wrapping of the phase. The plot shows phase in the range from $-\pi$ to $+\pi$. When the phase wraps the plot changes from pure white to pure black.

1.3.1 Limitations of paraxial ray tracing

The paraxial ray tracing results shown in this chapter were expensive to produce. A very robust ray tracing scheme provided by Thorbjorn Rekdal was used. An initial fan of rays is extrapolated from each source point, and the spatial separation of adjacent rays is then examined. If a pair of rays is more that 75m. apart anywhere in the model a new ray is interpolated between them. This continues until the coverage criterion has been met or until a minimum angular separation of 10^{-4} degrees has been reached. This scheme is slow but it ensures that there is reasonable ray coverage in the whole model. After these criteria have been met, the paraxial approximation is used to estimate traveltimes, amplitudes and phases at every point in the model. On average, six hundred rays were used in the final traveltime estimation at each shot, but many more rays were shot to obtain full spatial coverage. In comparison, the band-limited Green's function estimation used a constant polar mesh for each shot. The grid had 256 samples in angle and 326 samples in radius. The regular mesh makes interpolation of the traveltimes to rectangular coordinates a much simpler operation.

The results shown for paraxial ray tracing were not created using the true slowness model. They were created in a slowness model that had been smoothed over a radius of 100m. This is because the ray tracing is unstable in the true slowness model. The smoothing was chosen to be approximately one wavelength in length at the dominant frequency, and this was sufficient to stabilize the ray tracing. Ray-tracing in a model smoothed by less than 50m. was always unstable, smoothing over greater distances stabilized the ray tracing but gave results that were not good matches to the true wavefield, and did not produce good migration images.

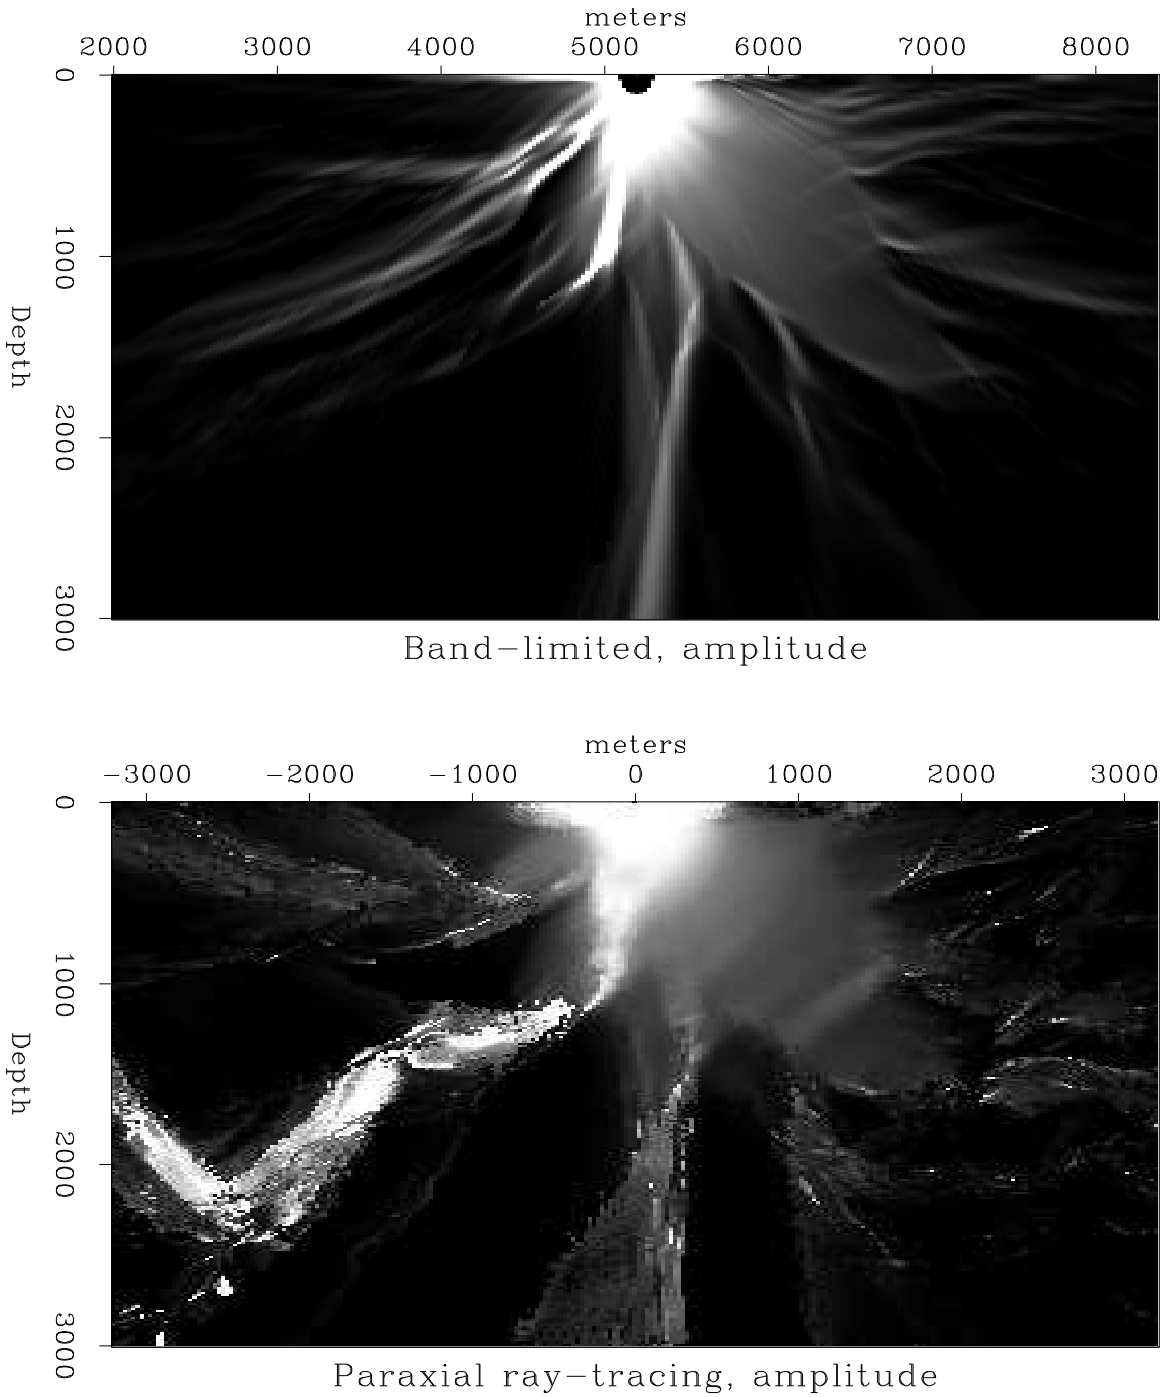

Figure 1.11: Comparison of amplitude fields from paraxial ray tracing and band-limited Green's functions. MarmTT-comp-amp [ER]

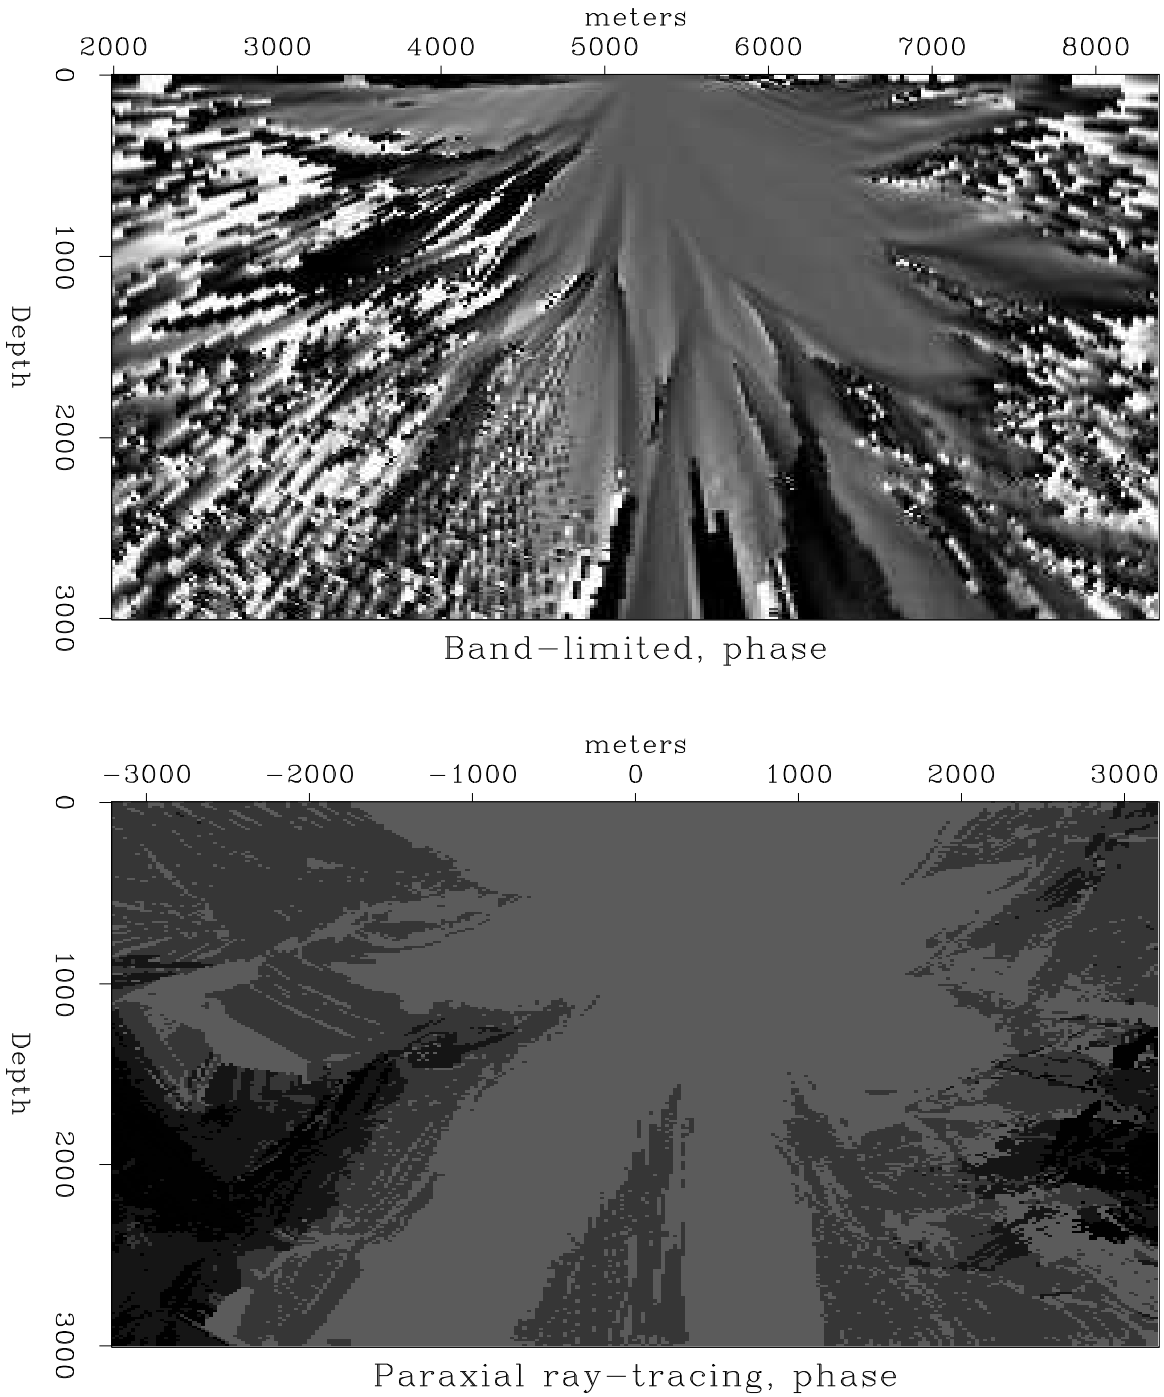

Figure 1.12: Comparison of phase fields from paraxial ray tracing and band-limited Green's functions. MarmTT-comp-phase [ER]

The smoothing can be interpreted as a method of partially accounting for the band-limited nature of the wavefield. A finite frequency wave is affected by some local average of the slowness field, not the slowness field at a single point. This concept is taken further in th work of Biondi (?) and Lomax (?). They each attempt to estimate an effective slowness field for ray tracing that is frequency dependent. The band-limited Green's function is calculated in the true slowness model because the wave equation is stable for all models. The smoothing effect is contained in the wave equation and does not need to be approximated. In the next chapter I show that the images created using band-limited Green's functions are a better match to full wavefield extrapolation than those created using paraxial ray tracing Green's function.