Detection of reservoir quality using Bayesian seismic inversion

James Gunning
CSIRO Petroleum, Ian Wark Laboratory,
Bayview Ave., Clayton, Vic. Australia,
Michael E. Glinsky
BHP Billiton Petroleum
1360 Post Oak Blvd. Houston, TX, U.S.A

ABSTRACT

We show how to invert seismic data for a permeable rock *sorting* parameter by coding the floating grain model of DeMartini and Glinsky (2006) into the Bayesian seismic inversion code *Delivery* (Gunning and Glinsky, 2004). The Bayesian prior embeds the coupling between elastic properties, porosity, and the floating–grain sorting parameter. The inversion uses likelihoods based on seismic amplitudes and a forward convolutional model to generate a posterior distribution containing refined estimates of the floating grain parameter and its uncertainty. The posterior distribution is computed using Markov Chain Monte Carlo methods. The test cases we examine show that significant information about both sorting characteristics and porosity is available from this inversion, even in difficult cases where the contrasts with the bounding lithologies are not strong, provided the signal to noise ratio of the data is favourable. This holds true even in the more difficult test case we examine, where the laminated reservoir net–to–gross is not significantly improved by the inversion process.

INTRODUCTION

Seismic data have long been highly valued as the most important information in delineating reservoir architecture and overall hydrocarbon–in–place in the oil exploration business. This is especially the case in regions where 'soft rock' characteristics make the presence of hydrocarbons visible in reflected amplitudes. If source–rock and charge interpretations are favourable, an attractive hydrocarbon–volume estimate from seismic amplitudes makes a compelling case for further appraisal work, such as drilling more appraisal wells.

But a commercial reservoir needs much more than favourable reserves – at the very least, the lithologies present must have favourable permeabilities for a commercial development to be viable. The value of seismic data in inferring permeability has been much more questionable, however, since flow characteristics of rocks are usually weakly coupled to their acoustic behaviour in well sorted rocks. The folk explanation of the poor coupling is that, loosely speaking, permeabilities are controlled by grain size, whereas acoustic properties are controlled by porosity, which is roughly independent of grain size for random packings of uniform size grains.

Such at least is the received wisdom for common well sorted sedimentary rocks. For poorly sorted rocks, less confidence can be placed in this generic belief. If they are known to be present, it is crucial to be able to distinguish their effects from that of clay, so a vital piece of contextual knowledge is a credible depositional story for the distribution or absence of clay, plus the spatial location of the poorer sorted rocks in the well data. It is also clear that modelling of poorly sorted rocks must be well informed by thin–section analysis, or an experimental equivalent. Once this context is established, the velocity relationships evident in the wireline log–data analysis must be explained by a consistent model.

In most established exploration or production basins in the world, oil companies maintain a database of ‘loading trends’ (the mutual dependencies of elastic properties, density and porosity on burial depth or effective stress) applicable to the reservoir rocks in the region. These are usually the better quality or ‘clean’ reservoir rocks. Challenges arise when new log or core data appears which deviates substantially from these trends. In the simplest case, such deviations can be attributed to the presence of clay, either laminated or dispersed, and the analysis can proceed with a conventional clay model. But there are cases when anomalous velocity behaviour in the log and core measurements is not explicable using conventional clay theories, and more complex effects due to poor sorting constitute a better explanation of the data. Further, if the anomalous velocity measurements of this kind are well correlated with permeability changes, a theory couched in terms of sorting variations may well be more suitable than one based on clay.

As an example, if there is no evidence of dispersed clay, one can imagine poorly sorted rocks in which fine small grains may simultaneously have the effect of constricting the fluid flow, as well as increasing the density without providing compensatory load support, which will correlate the seismic velocity with the permeability. Such a model is developed in DeMartini and Glinsky (2006). We use this model as the basis of the inversion studies presented in this paper.

Settling the issue of the appropriate rock physics model to describe the observed log and core data is only the first problem. The second is the question of what may be sensibly and defendably inferred from seismic amplitude data given the parameters of the regional rock–physics model. This question of inferability must be settled by a careful model which incorporates a reasonable level of uncertainty in all the other factors which affect seismic amplitudes. Simple ‘maximum–likelihood’ inversion studies are only a part of the answer: the uncertainty in the inversion is the more key ingredient. For example, in the absence of interference or ‘tuning’ effects, migrated seismic amplitudes are sensitive to the velocities, densities and fluid content of the rocks both above and below a lithological boundary. These will vary with depth and effective stress. A key question is then the extent to which changes in the sorting characteristics of the reservoir rock are masked by natural variations in the other factors, the answer to which is crucial in interpreting seismic amplitudes.

To address these issues rigorously, we have coded the ‘floating grain’ model presented in this paper into the Bayesian seismic inversion tool *Delivery* (Gunning and Glinsky, 2004; Gunning, 2003) which properly accounts for all the requisite uncertainties in the rock physics and stratigraphic geometry in realistic models. The suite of stochastic ‘posterior’ models produced by *Delivery* then determines what degree of refined knowledge about the sorting characteristics of the rock properties is available from seismic data.

For some geological environments, we believe that the presence of poorly sorted material, which is strongly correlated with permeability in core tests, has a detectable influence on acoustic properties, and is thus inferable from seismic data. In this paper we present two such forward modelling studies to validate this belief, one for the simplest 'clean interface' problem of a shale seal bounding the main reservoir sand, and the second problem is a representative multi–layer inverse problem typical of the depositional architecture of the prospect whose data set inspired this study.

The layout of this paper is as follows. The section *Rock physics models* summarises the floating grain model and the methods used to estimate its regression parameters from the data of interest. Details of how this is coerced into a form suitable for *Delivery* follow. There are several implementation details relevant to the way in which *Delivery* treats the prior which are important, but we defer these to Appendix B. In *Numerical Examples* we illustrate the implications of this model for two inversion problems – the first (*Example A: A simple model system*) being the simplest seal over reservoir toy problem, where we explore the characteristics of the prior in some detail. In *Example B: More complex model based on field data*, a much more fully fledged model incorporating a complex seal structure and two reservoirs is developed. We summarise our findings in the *Conclusions*.

ROCK PHYSICS MODELS

We recount briefly the rock–physics model described in DeMartini and Glinsky (2006). It assumes measurements apply in the Gassmann low–frequency limit, and that the reservoir is a homogeneous isotropic medium. In general we distinguish between the fluid porosity ϕ and structural porosity ϕ_s. If the grain density and bulk modulus are ρ_g and K_g respectively, then filling the pore space with a fluid or suspension with properties ρ_f, K_f produces an effective medium of density

$$\rho = \rho_g(1-\phi) + \rho_f\phi \tag{1}$$

and compressional and shear velocities

$$v_p^2 = \frac{K_g}{\rho}\left(\frac{3(1-\nu_m)}{1+\nu_m}\beta + \frac{(1-\beta)^2}{\phi_s(K_g/K_f - 1) + 1 - \beta}\right), \tag{2}$$

$$v_s^2 = \frac{K_g}{\rho}\frac{3(1-2\nu_m)}{2(1+\nu_m)}\beta. \tag{3}$$

Here the dimensionless matrix bulk modulus $\beta = K_m/K_g$, and ν_m is the matrix Poisson's ratio. The dependence of the matrix β on (structural) porosity is taken to be that of a conventional critical–porosity model (Mavko et al., 1998; Nur et al., 1991)

$$\beta(\phi_s) = (1 - \phi_s/\phi_c)^\lambda. \tag{4}$$

where ϕ_c is a critical 'suspension' porosity, usually around 0.42, and λ a data–fitted constant.

For a poorly sorted collection of grains, the 'finer' grains are treated as a secondary component, which contributes in two pieces: (i) some small volume fraction of fine grains do not support the rock matrix and act like a pore-space fluid, while (ii) the remaining fraction is bound or captured into the load–bearing frame as the rock is buried over time. If the overall fraction of small grains introduced is f_*, and a fraction f_c of these are 'captured', then $\phi_{\text{flt}} \equiv (1-f_c)f_*$ is the volume fraction of 'floating' grains. This floating fraction is treated as an effective fluid and modelled via Gassman substitution, while the effect of the captured grains is absorbed by using the structural porosity $\phi_s = \phi + \phi_{\text{flt}}$ in β, equation 4, and the load–bearing or structural porosity appearing in the denominator of 2. The effect of the bound grains on the matrix Poisson's ratio can be shown to be weak, and is thus neglected.

In DeMartini and Glinsky (2006), extended arguments are presented as to why treating the second component as a significantly different mineral does not yield an adequate match to the experimental data shown below. Core and thin–section analysis also precludes the presence of dispersed or laminated clay. The data is then modelled using a 'bi–modal' (large, and small floating grains) mixture of a single mineral, which has the advantage that the overall $p-$wave velocity simplifies to the expression

$$v_p^2 = \frac{K_g}{\rho_g(1-\phi)+\rho_f\phi}\left(\frac{3(1-\nu_m)}{1+\nu_m}\beta(\phi+\phi_{\text{flt}}) + \frac{(1-\beta(\phi+\phi_{\text{flt}}))^2}{\phi(\frac{K_g}{K_f}-1)+1-\beta(\phi+\phi_{\text{flt}})}\right). \quad (5)$$

Arguments are also furnished to demonstrate that the v_s vs v_p regional trend does not change to leading order under this model, a prediction which is corroborated by the data.

The actual $\{v_{p,i}, \phi_i\}$ data from log measurements in the province of interest appears to fall into distinct clusters. One cluster, associated with core data of decent permeability and fairly monodisperse sands, is modelled as 'clean' rock, with no floating component ($\phi_{\text{flt}} = 0$). This data is used in a non–linear regression $v_i = v_p(\phi_i, \phi_{\text{flt}} = 0, \lambda) + \epsilon_{p,i}$ for the exponent λ, assuming generic values for quartz ($K_g = 37$Gpa, $\rho_g = 2.65\text{g/cm}^3$), brine ($\rho_f = 1.05\text{g/cm}^3$, $v_{p,f} = 6.07$km/s), a mid–porosity range typical clean sand Poisson's ratio ($\nu_m = 0.15$), and a critical porosity $\phi_c = 0.42$. The regressed fit of $\hat{\lambda} = 1.724$ is then used in a second regression for ϕ_{flt}, using only the anomalous data, yielding $\hat{\phi}_{\text{flt}} = 0.039$, which is a plausible average value for the anomalous data. The data and fits are shown in Figure 1. Note that the 'anomalous' data is probably a mixture of rocks with variable sorting characteristics, ranging from nearly clean to perhaps $\phi_{\text{flt}} = 5\%$.

Since the floating fraction is an unobserved quantity for each measurement, an important statistical question is how its distribution can be disentangled from the regression variance ϵ_p. The underlying distribution of ϕ_{flt} is unknown, but is likely to contain clusters, often coinciding with well–groups. We have chosen to fix the variance of ϵ_p to that of the regional trend, which accounts for the dispersion in the velocities due to 'conventional' effects, and attribute the remaining dispersion to the effect of ϕ_{flt}. A multi-cluster analysis (using a modified form of MCLUST (Fraley and Raftery, 2003)) of the univariate distribution of the residuals $\xi_i = v_{p,i} - v_{\text{reg.trend}}(\phi_i)$, with cluster 1 fixed to the regional error (mean 0, variance $var(\xi) = \sigma^2_{v_p,reg.}$) shows the most statistically significant clustering model is a two cluster split. The data points grouped with the regional trend are shown circled in Figure 1.

Linearisation of the best fit velocity relation in ϕ, ϕ_{flt} about a suitable mean porosity $\bar{\phi}$ and $\phi_{\text{flt}} = 0$ yields a straight line fit

$$v_p = a_p + b_p(\phi_{\text{flt}} + (1+g)\phi) + \epsilon_p, \quad (6)$$

where we have written the three required constants in this way for consistency with the notation of DeMartini and Glinsky (2006). The (zero mean) error term is ϵ_p. The linearisation is clearly reasonable for the data clusters of Figure 1.

In DeMartini and Glinsky (2006), the response to loading is captured using a standard exponential regression model dependent on the effective stress σ_{eff}, with an additional term

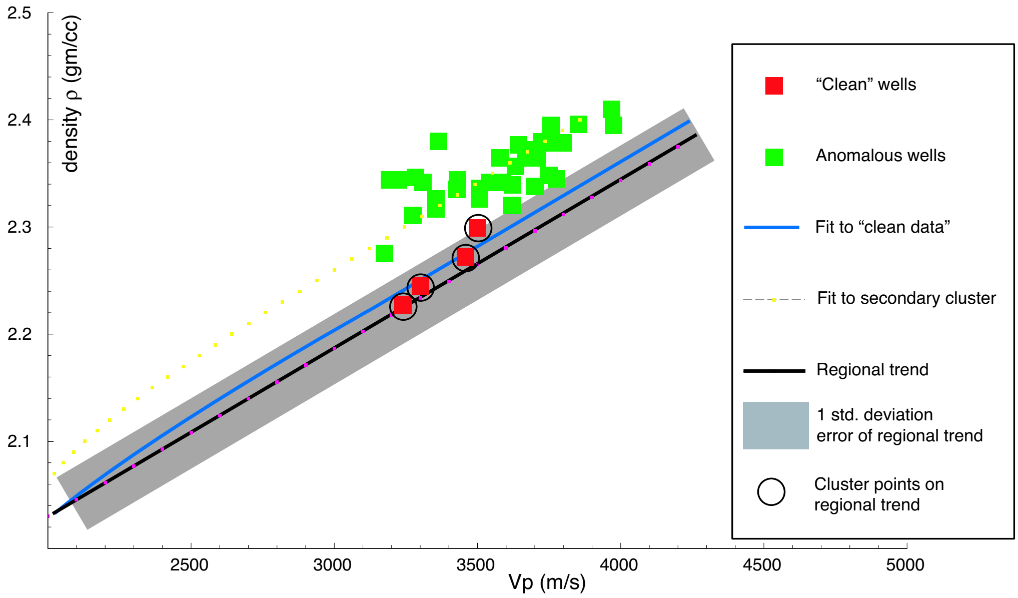

Figure 1: Density vs. v_p data for 'clean' sands and anomalous sands, together with non–linear regression best fits ($\lambda = 1.724$ and $\phi_{\text{flt}} = 0, 0.039$ respectively). The straight line 'regional trend' is the average across a much larger data set for the region, with approximate errors as shown.

describing the grain capturing effect:

$$\phi = a_\phi + b_\phi(1 - e^{-\sigma_{\text{eff}}/P_0}) - \frac{\phi_{\text{flt}}}{1 - f_c} + \epsilon_\phi. \tag{7}$$

This expression is a statement that the total space occupied by the pore fluid and the finer grains is compressed under loading in a 'conventional' way, and amounts to a definition of the capture fraction f_c. Basin modelling provides estimates of the effective stress, and the floating fraction estimates from the velocity regression on data clusters then provide a way to estimate a_ϕ and f_c through fitting. When this process is performed on the data set here, estimates of $f_c = 0.3522$, $a_\phi = 1.1$, $b_\phi = -0.8759$, $P_0 = 800\text{PSI}$ and $\sigma_{\epsilon_\phi} = 0.0024$ are produced.

Conversion to *Delivery* formats

The *Delivery* software has an established style of representing loading or compaction curves where p-velocity is regressed directly against suitable loading terms (Gunning and Glinsky, 2004). The fully linear form of the prior and the assumption of Gaussian regression errors also enables a multi–Gaussian prior to be formulated. The naturally 'augmented' version of the *Delivery* regressions suitable for incorporating sorting effects is thus

$$v_p = A_{v_p} + B_{v_p} d + C_{v_p} \text{LFIV} + D_{v_p} \phi_{\text{flt}} + \epsilon'_p \tag{8}$$
$$\phi = A_\phi + B_\phi v_p + C_\phi \phi_{\text{flt}} + \epsilon'_\phi, \tag{9}$$

using the existing notation in Gunning and Glinsky (2004). The shear relation is unchanged. The prior for each reservoir layer thus has the conditional form $P(v_s|v_p)P(\phi|v_p, \phi_{\text{flt}})P(v_p|d, \text{LFIV}, \phi_{\text{flt}})$, where the variance of each of these Gaussian components comes from the regression variance, e.g. $\mathrm{var}(\epsilon'_p) = \sigma_p^2$. The effective–stress dependence can be effected by taking the 'loading depth' $d \equiv (1 - \exp(-\sigma_{\mathrm{eff}}/P_0)))$ if required (we use the `depth_rock_curves` entry in the xml file used by *Delivery*, to distinguish from model depth). The LFIV term is then unnecessary, and is dropped by setting $C_{v_p} = 0$. Approximate conversion to this *Delivery* style of the coupled regressions can be obtained with the formulae of Appendix B.

NUMERICAL EXAMPLES

We present here two examples illustrating how the inversion works using a floating grain model fitted to data from the province of interest. The first example is the standard test problem of a single isolated reflector, and the second a more complex model based on the full log data. The region of interest contains various lithologies, but the main cap rock above the oil–bearing sand is a shale, so we are chiefly interested in the trend properties of the pay sand and this overlying shale. The shale and sand trends are common to both models, so we dispense with these first.

Rock trends for sands and shales

Shales

Standard shale trends for the area in question are, with z as depth below mudline:

$$\begin{aligned} v_p &= -5377.26 + 0.9457z \pm 476 \text{ (ft/s)} & (10)\\ \rho &= 0.5343 v_p^{0.166} \pm 0.029 \text{ (gm/cc)} & (11)\\ v_s &= -3373 + 0.8012 v_p \pm 206 \text{ (ft/s)}. & (12) \end{aligned}$$

Sands

For the v_p relation 5, the 'clean trend' applicable is obtained by a linearisation of the fit line shown in Figure 1 (the nonlinear fit to the 'clean rock' data cluster), whose maximum–likelihood fit has $\lambda = 1.724$ and material constants as per the accompanying description. At the mean data porosity $\bar{\phi} = 0.24$, the linearisation of equation 5 to the form 6 yields constants $a_p = 18795.8$, $b_p = -30970.6$ and $g = 0.0247936$, with estimated error $\sigma_p = \sqrt{\mathrm{var}(\epsilon_p)} = 750$ ft/s.

The loading trend for the pay sands is established in DeMartini and Glinsky (2006) as

$$\phi = 1.1 - 0.8759 Z_{\mathrm{eff}} - 1.5437 \phi_{\mathrm{flt}} + \epsilon_\phi. \qquad (13)$$

with $Z_{\mathrm{eff}} \equiv (1 - e^{-\sigma_{\mathrm{eff}}/800\mathrm{PSI}})$ and error $\sigma_\phi = \sqrt{\mathrm{var}(\epsilon_\phi)} = 0.0024$. Since the shale trends are against z, we have converted this stress regression 13 to a depth trend, since the loading term Z_{eff} is very nearly linear in depth over the depth range of interest. The equivalent pay sand depth trend (c.f. equation 7) is

$$\phi = 0.525566 - 1.59914 \times 10^{-5} z - 1.5437 \phi_{\mathrm{flt}} + \epsilon_\phi. \qquad (14)$$

The last regression coefficient ($1.5437 = 1/(1 - f_c)$) corresponds to the capture fraction $f_c = 0.3522$. With the understanding that z now plays the role of $Z_{\rm eff}$, the conversion formulae of Appendix B yields the *Delivery* style constants

$$\begin{aligned} A_{v_p} &= 2115.21 \\ B_{v_p} &= 0.507541 \\ D_{v_p} &= 18025 \\ A_\phi &= 0.592251 \\ B_\phi &= -3.151 \times 10^{-5} \\ C_\phi &= -0.97581 \end{aligned} \qquad (15)$$

The shear trend for the data regressed directly to yield

$$v_s = -3996 + 0.8940 v_p \pm 226 \text{ (ft/s)}. \qquad (16)$$

We have used the clean–sand regional–trend error estimates $\sigma'_{\epsilon_\phi} = 0.0093$ and $\sigma'_{\epsilon_p} \approx$ 344ft/s in the augmented model, corresponding to the 'conventional effects' assumption described earlier.

Example A: A simple model system

The aim here is to determine whether the presence of floating grain material in the reservoir rock is ascertainable from reflected amplitudes using the rock physics model and regional regressions just derived. We begin with the simplest 2–layer shale/sand system, which is free of the complication of interference or tuning effects.

The well logs used to construct the prior have some clean rocks (no floating grains) and rocks with floating content of around 2–5%. To model the inferability of the float fraction, we constructed synthetic seismic truth–case stacks for near and far stack angles of a few degrees and about 30°, using a truth case model with 5% float and NG = 1, and all other parameters at the most likely values from the trends. The reservoir fluid is taken as brine for this simple study. Figure 2 illustrates the system, with truth-case plus posterior near and far synthetics from the posterior of case iii) we describe shortly.

For inversion, the prior on floating fraction is taken as $N(0, 0.05^2)$, and we attempt to compute the posterior floating fraction from three cases: i) net–to–gross (NG) fixed at 1, near–stack only, ii) (NG $\sim N(1, 0.2^2)$, near–stack only, and iii) (NG $\sim N(1, 0.2^2)$, near and far stacks.

For case i), Figures 3– 5 show 3 possible forms of the prior – varied for illustrative purposes – on which we superpose the original log data, which illustrates how the floating grain effect smears out the regional prior.

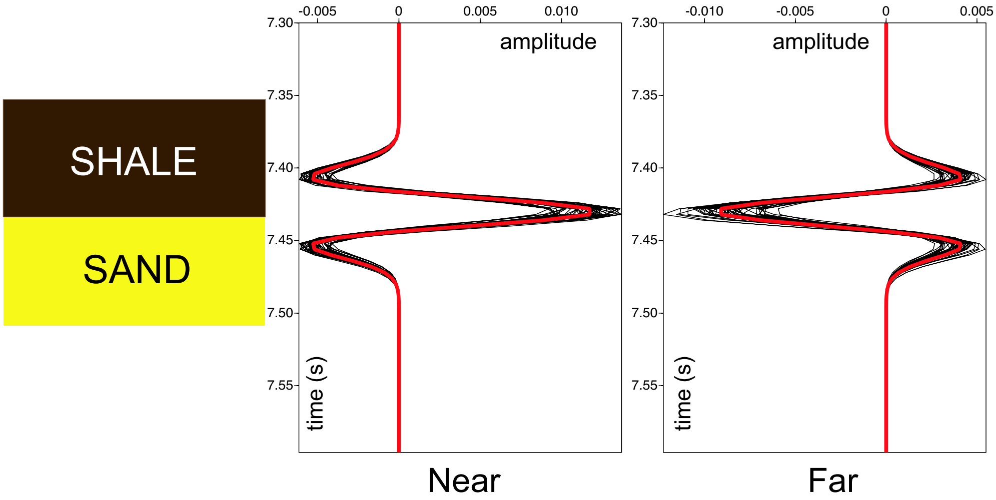

Figure 2: Two–layer model system with truth case (thick red) seismic traces and synthetics from the posterior (black) for the inversion case iii) described in the text. The absolute noise level is set at 0.002 for both stacks.

Inversion Analysis of Posteriors

Fixed NG, single (near) stack

Figure 6 shows salient scatterplots of properties of the sand layer before and after inversion. The peak–signal to noise ratio is set at about 6:1. This and subsequent figures use symbols defined in Gunning and Glinsky (2004) as follows: i) R_near and R_far, defined as R_{pp} for the near and far stack, from equation 22, ii) overall layer effective density $\rho_{\rm eff}$ and velocity $v_{\rm p,eff}$, defined in equations 9 and 10. The inversion is clearly able to detect the presence of floating grain material and refine the porosity estimates.

Free NG, single (near) stack

Figure 7 shows salient scatterplots of properties of the sand layer after inversion, where the model has additional net–to–gross freedom in the prior NG $\sim N(1, 0.2^2)$. As shown inset in Figure 7.D, the inversion produces virtually no posterior sharpening of the net-to-gross distribution. Nevertheless, the floating grain fraction estimate is still noticeably improved.

Fixed NG, near+far stack

Inversion using shear data in principle may help narrow down floating grain porosities better, as the shear carries additional information. The far stack for this test case is set at about 30 degrees (c.f. a few degrees for the near) and the reflected amplitude is much weaker (AVO effects). The noise level was set at the same value as for the far stack.

For this case, it turns out that relatively little improvement in the estimates of the primary quantities $\phi_{\rm flt}$ and porosity occurs. A significant sharpening in the shear velocities

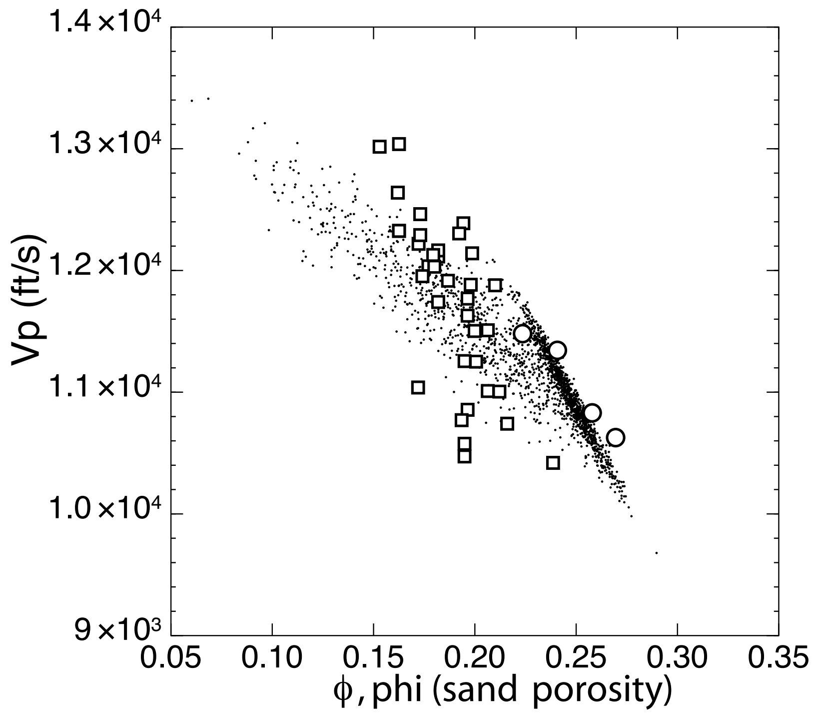

Figure 3: This and figs. 4 and 5 show three "pedagogic" priors for velocity *vs* porosity in a pure sand layer. Clean well data points (circles) and float-polluted well data points (squares) are plotted on all three graphs: dots ($\cdot$) are draws from the model prior. (a) Prior constructed with artificially narrow $\sigma_\phi = 0.002$, showing how parameters arise from a clear 50:50 mixture of 'clean–rock regression' points and an elliptical smear from the effects of the floating grains. The clean rock trend is obviously far too narrow to embrace the clean well (+) measurements, but the 'clean trend' is clearly visible and centred. Recall the well data is from a spread of depths, so it is not expected that the prior (applicable at the reservoir depth only) covers all well data.

v_s and especially the overall far–reflection coefficient R_far does occur. But the coupling of these quantities with the primary quantities of interest induced by the prior and likelihood terms does not appear to be sufficient to significantly improve their estimates, in this particular case.

In summary, the test case here appears to show that an inversion is capable of detecting the presence or absence of floating grain material for the kinds of rocks studied at the depths of interest, as well as 'tune up' the reservoir porosity estimate. The basic reflection coefficients are relatively weak ($|R| \sim 0.03$) for both near and far stacks at this depth, so it is clear apriori that the problem will be difficult. In contrast to our usual experience of Bayesian inversion with imaged seismic data, the updates to the sand/shale net–to–gross are very weak, despite the encouraging results for floating–grain fraction. This phenomenon is a particular conspiracy of the impedance trends for the rocks in question, so this asset appears to be a particularly challenging example. Further challenges arise in the more complex field example of the next section.

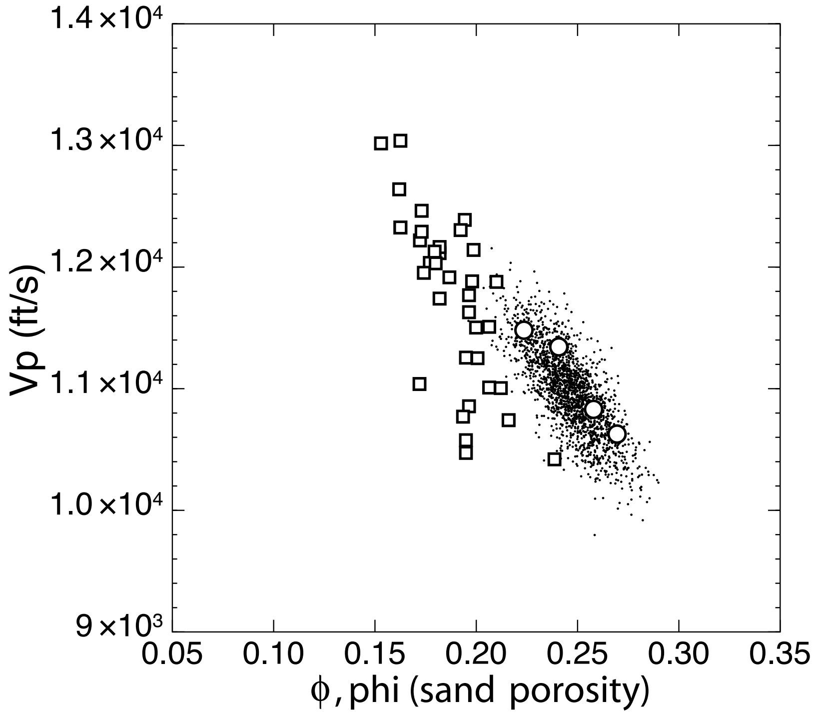

Figure 4: Illustrative prior (b): prior drawn from clean rocks only (float–fraction $\phi_{\rm flt} \sim N(0,0)$), with broad porosity uncertainty $\sigma_\phi = 0.02$. The tails of the distribution do not contain the floating grain data comfortably. Clean well data points (circles), float-polluted well data points (squares); dots ($\cdot$) are draws from model prior.

Example B: More complex model based on field data

In this example, based on actual field data, the complicating effects of additional lithologies and tuning considerations come into play. The oil reservoir system we model here features 'upper' and 'lower' pay sands which are capped by a complex draping structure including thick, acoustically hard marls and thin, soft, silty layers. An overall simplified 6–layer sequence for the system has been modelled as {(1) Marl, (2) silt-marl-stringer-complex, (3) shale, (4) 'upper' sand, (5) shale, (6) 'lower' sand}, where the silt-marl-stringer-complex is an upscaled (impermeable) layer absorbing some of the very thin structures in the cap. The near offset reflectivity from the marl edge is sufficiently strong to dominate the reflection from the 'upper' sand top, so much extra information comes from the interface with the shale below, and an additional 30° stack whose p-wave reflectivity is more comparable to the marl.

Figure 8 shows the upscaled layer structure superposed on the logs and shale plus sand trends for the model. Note there are distinct trends for the marl and stringer complex.

The main pay sands are not very clean, and are estimated to have net–to–gross values of around $65 \pm 10\%$. Oil is proved in these reserves with saturations of around 60%. For simplicity, we fix the oil probability (1) and saturation in these layers. The floating grain fraction in the main reservoirs is given a prior of $N(0.02, 0.03^2)$ (truncated at 0, naturally), which gives a significant prior probability to the zero–float or clean–sand case. The 'truth

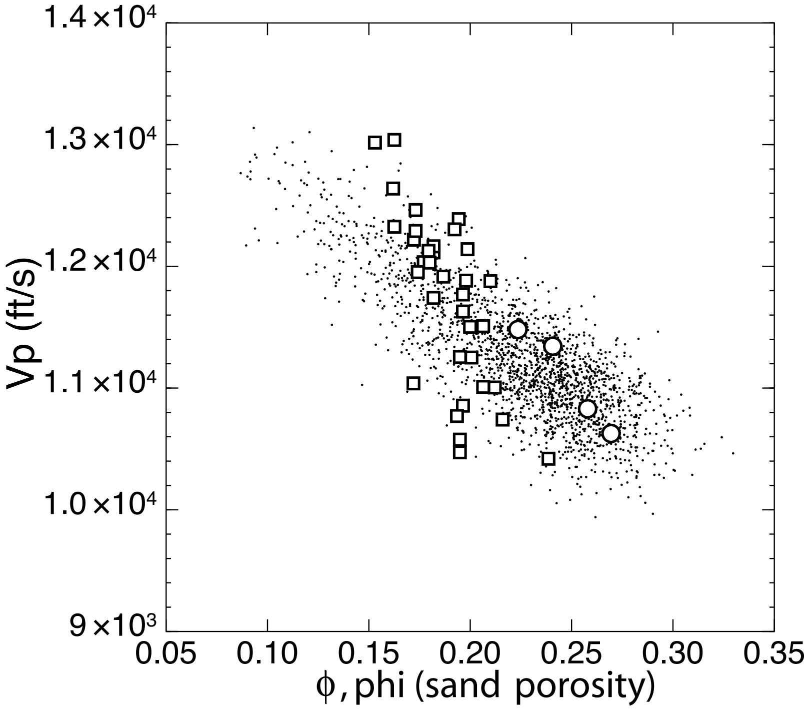

Figure 5: Illustrative prior (c): prior with $\sigma_\phi = 0.02$ and float–fraction $\phi_{\rm flt}$ distributed as $\sim N(0, 0.05^2)$. Only a few measurements appear to lie at the periphery of the distribution, but the mixture character is not as clear to the eye as in Figure 3. Clean well data points (circles), float-polluted well data points (squares); dots $(\cdot)$ are draws from model prior.

case' corresponds to a float fraction of 3.5%.

Figure 9 shows some typical 'spaghetti plots' of the synthetic seismic from the posterior plotted against the 'truth case' data, for both stacks. Typical realisations from prior and posterior are also shown as 'layer-cake' images of layers against realisation number.

As might be expected, the inversion produces strong updates to parameters like the layer times, impedances, effective layer velocity, and porosity. Some details are shown in figures 10 and 11. A reasonable improvement in the float-fraction estimate occurs, in particular the fact that the posterior significantly reduces the 'zero float' possibility. The most likely prediction is correct at around 0.035. The net–to-gross estimates are barely improved, mainly because the 'upper' sand offers a weak impedance contrast to its mixing shale when floating grain material is present at around 3% (c.f. the trend curves). The sensitivity to floating grain fraction is very much higher than that to NG, so the update is stronger. Again, the far–stack data helps to refine the shear properties significantly, but this does not couple back through the prior strongly enough to improve estimates of the quantities of direct interest markedly. Though not shown, the far stack data also markedly improves the v_p statistics of the shale in layer 5, but not the adjoining reservoir sands. Overall, however, the improvements over the prior are not strong in view of the aggressive signal to noise ratio.

In conclusion, the more realistic toy problem here shows that detection of poorly sorted

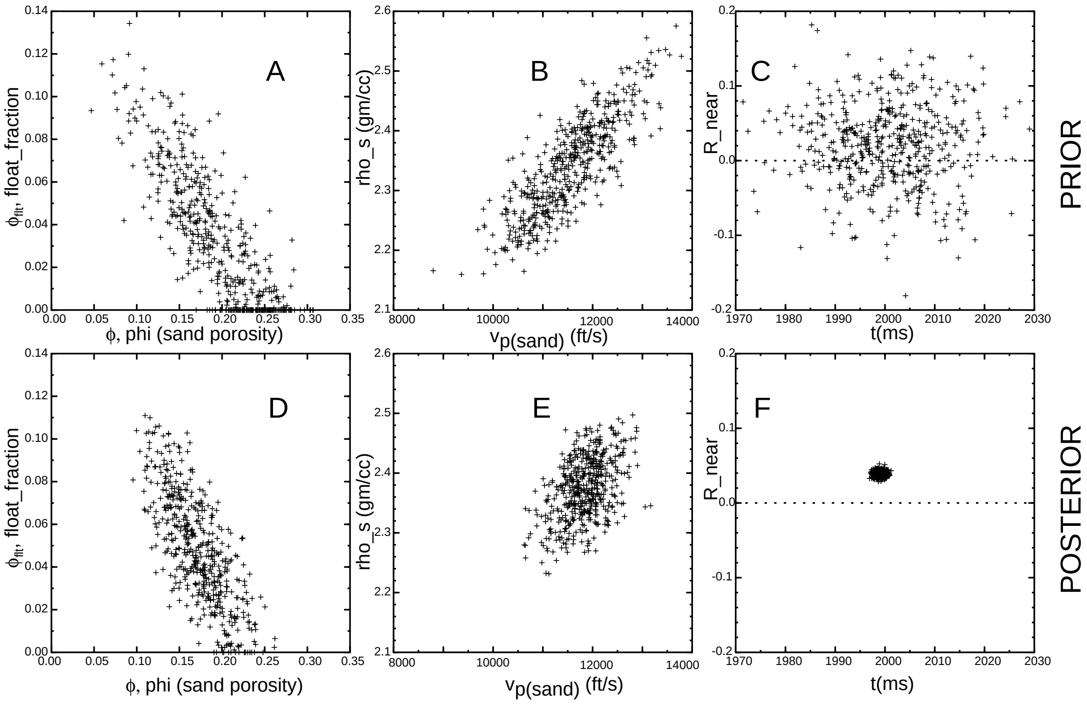

Figure 6: Scatterplots from prior distribution (top row) and posterior (bottom row) for 2–layer model with sand below shale. (A,D) float_fraction (ϕ_{flt}) vs sand porosity. A perceptible narrowing around the true answer of $\phi_{\text{flt}} = 5\%$ is visible: fewer clean sands are produced in the posterior, and (B,E) the sand density *vs* sand velocity $v_{\text{p(sand)}}$ scatterplot narrows more obviously. (C,F) The effective reflection coefficient R_{near} *vs* layer–time t is clearly pinned down sharply. As usual, these parameters are most heavily constrained.

material is possible with very favourable signal–to–noise ratios. Several asset–specific issues make this more difficult than might at first be expected. The first is that the reflection at the top of the 'upper' pay falls in the sidelobe of the very strong reflection from the overlying marl layer, so tuning effects and uncertainties in the modelling of the cap rock package in general limit what may be discerned about the underlying sand. Secondly, the particular loading behaviour characteristic of the pay depth makes the shale impedance quite close to that of the sand, so the overall strength of the main reservoir reflection is notably weaker than other nearby events. In view of the difficulties, it is relatively consoling that positive information about the floating grain contribution can be drawn when very little can be said about the net–to–gross.

CONCLUSIONS

The quantitative floating–grain rock physics model presented in DeMartini and Glinsky (2006) has been incorporated into the Bayesian seismic inversion program *Delivery*. Development of the trend models for use in *Delivery* requires careful log and core analysis and some simple nonlinear regression studies. The simple synthetic inversion studies we present are closely based on actual asset data, and show that genuinely improved estimates of the floating–grain or sorting characteristics, plus the reservoir porosity, are possible if the

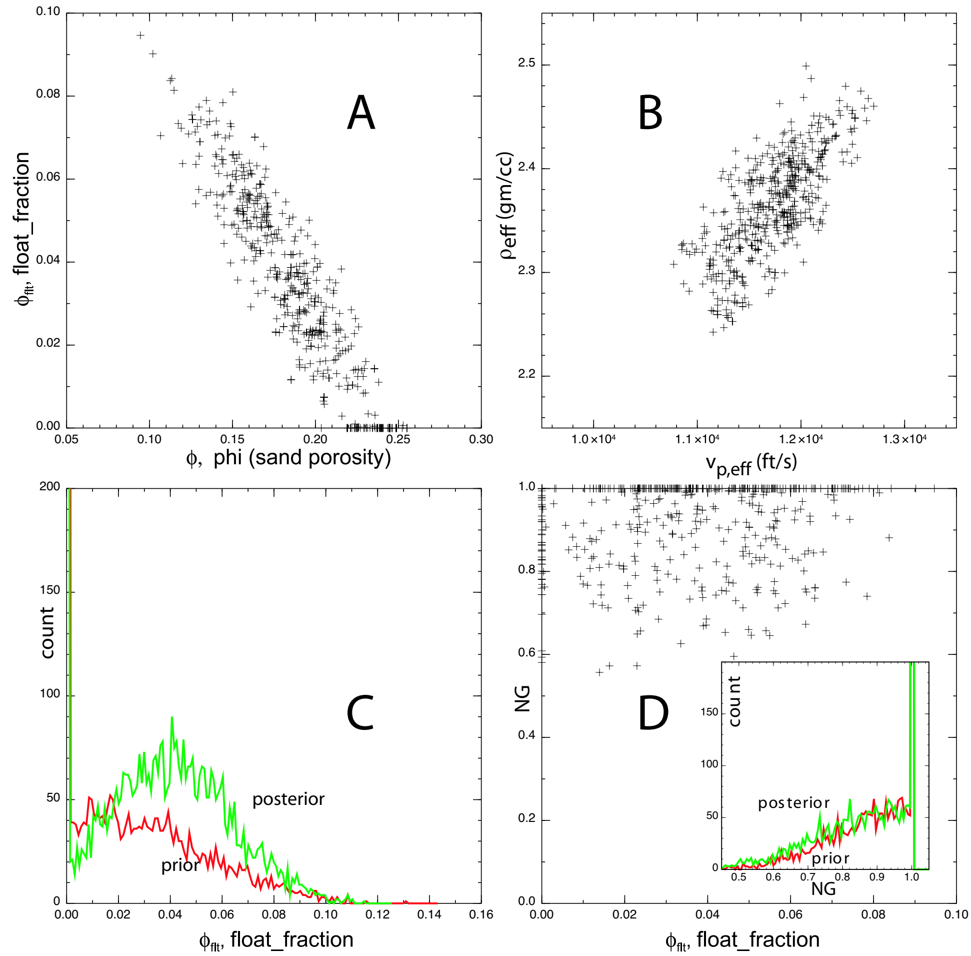

Figure 7: Scatterplots from the posterior for the model with looser net–to–gross (NG) distribution in prior. Inset (A) floating–grain fraction $\phi_{\rm flt}$ vs sand porosity ϕ, (B) effective density $\rho_{\rm eff}$ vs effective velocity ($v_{\rm p,eff}$). (C) Histogram of $\phi_{\rm flt}$ from the prior and posterior. (D) Scatterplot of NG vs $\phi_{\rm flt}$, with (inset) histograms of NG from the prior and posterior. Again, no obvious strong correlation appears here, with the density strongest near the truth case values $(0.05, 1.0)$.

seismic data has sufficient signal to noise ratio. For this particular study, far–offset data does not appear to provide markedly better updates of anything other than shear–related quantities.

The asset in question has some particular challenges associated with a relatively weak reservoir reflection coefficient and complex tuning interference from a hard marl above the reservoir cap, so we regard it as a difficult test case. The inversion techniques for rock quality now available via the floating grain model capability of *Delivery* can be expected to produce markedly sharper posterior updates for data sets free of these particular conspiracies.

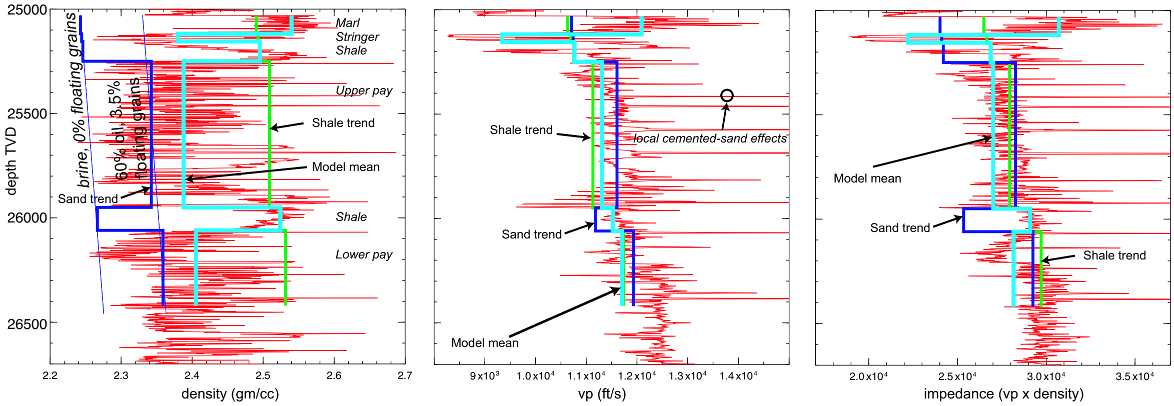

Figure 8: Layer based model truth case properties for the test model: logs (red) and truth-case mean properties (cyan) are shown, for density, p–wave velocity and normal impedance. The succession of 6 layers is clearly evident. The shale trend is shown green, sand blue: the deflections in the sand trend lines are due to the floating grain term, and these are shown without fluid substitution. The streaks in the logs are due to small cemented sand units which have no large scale effect on the seismic.

ACKNOWLEDGMENTS

The first author gratefully acknowledges generous funding from the BHP Billiton technology program.

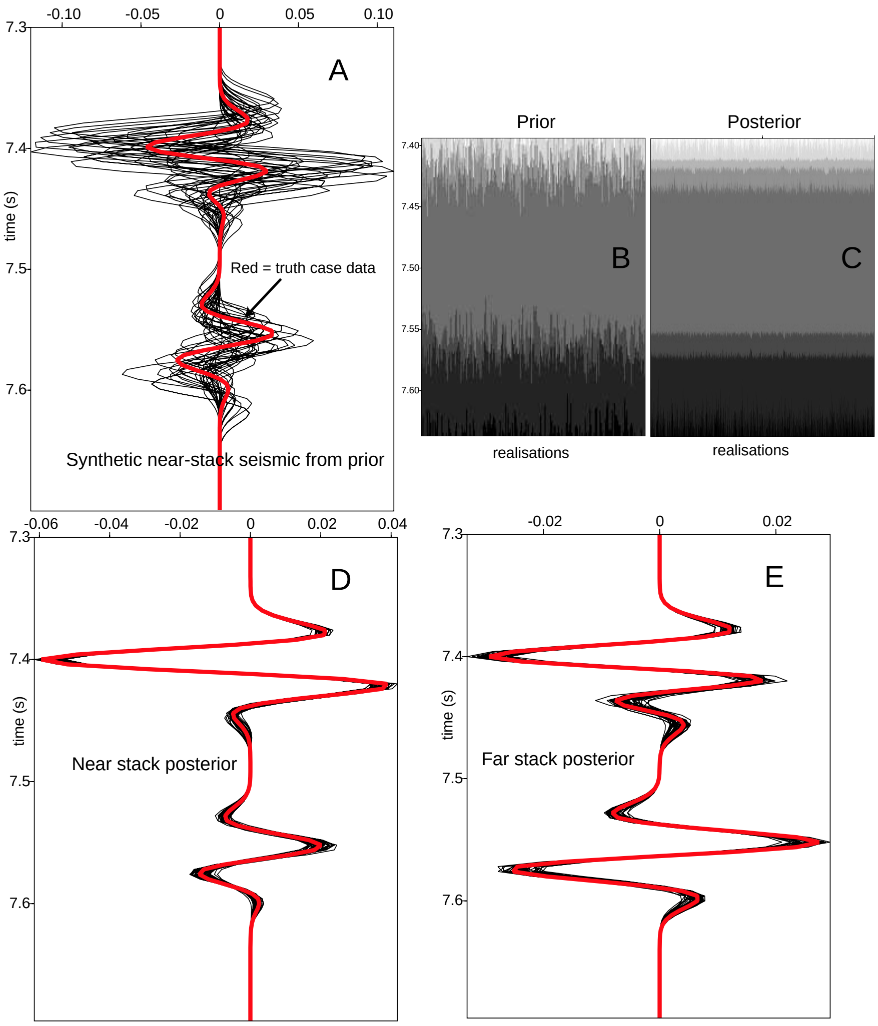

Figure 9: Insets A,D,E: Synthetic seismic curves (black) for model samples, superposed on actual data (red), as drawn from prior (A – near stack), and posterior (D – near stack; E – far stack). Note the noise level is set very low; the S/N is at least 10:1. Insets B,C show approx. 50 layer realisations from the prior and posterior respectively, in time. The layers are shaded in ascending grayscale according to layer number.

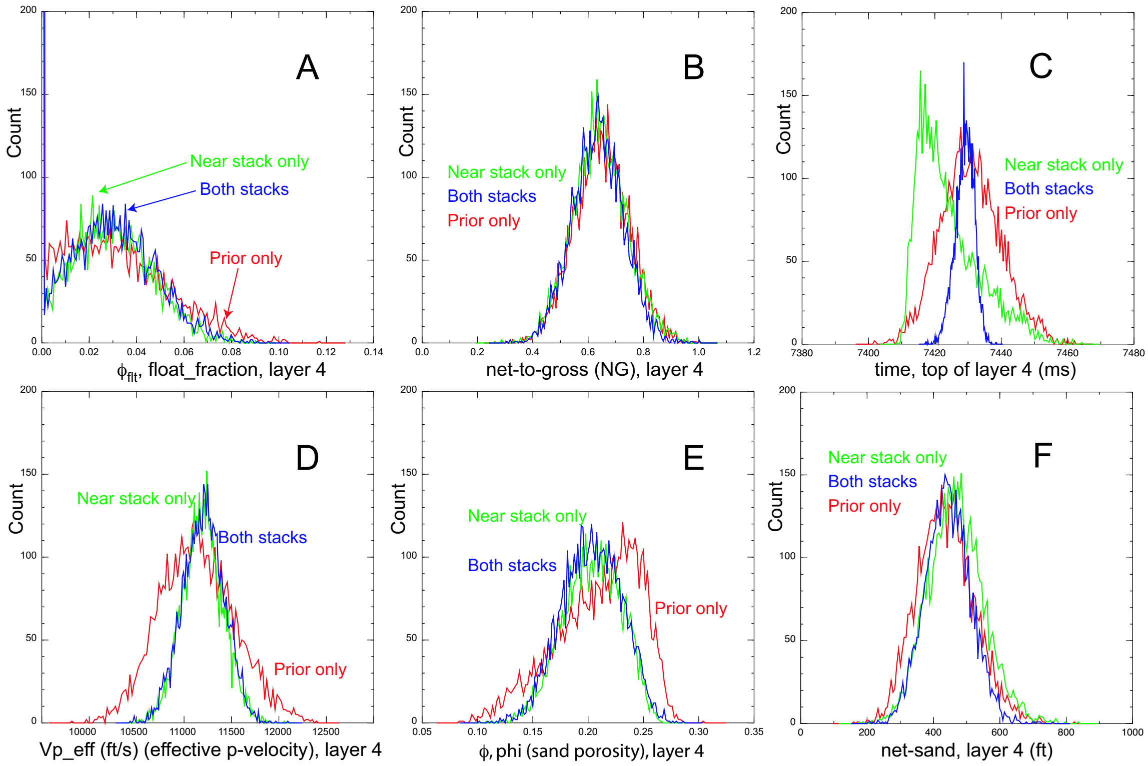

Figure 10: Prior (red), near-stack-only posterior (green), and both-stack posterior (blue) histograms for selected properties of the 'upper' sand (layer 4). The strange asymmetric layer–4 time posterior in C for near–stack only (green) is evidence that the top layering structure is not well resolved from this stack alone. More revealing detail appears in the scatterplots of Figure 11.

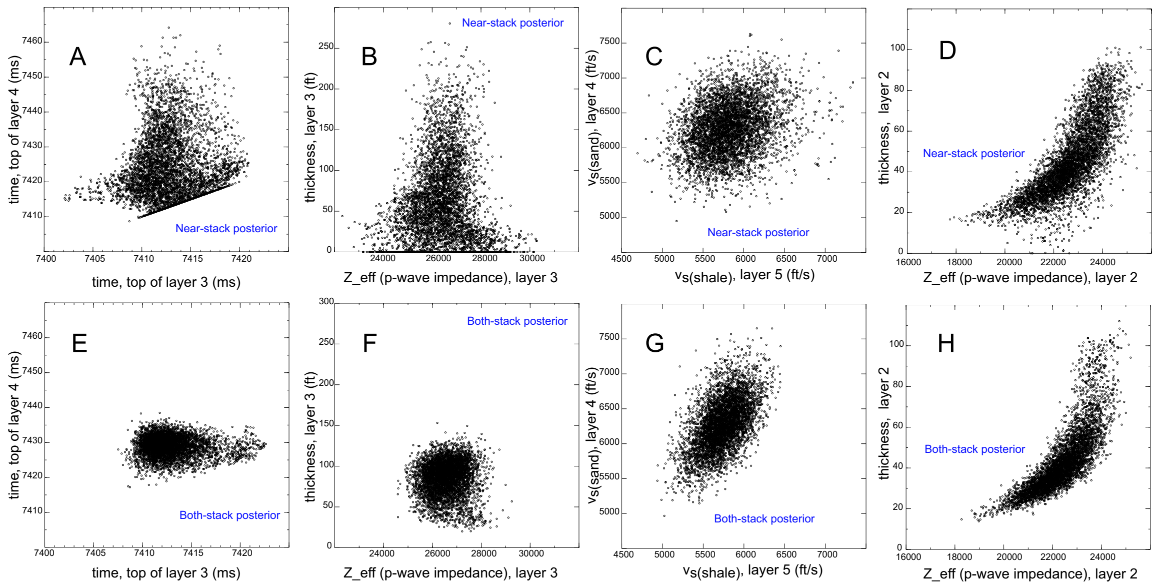

Figure 11: Insets A,B,C,D: Scatterplots of selected property–pairs from the near–stack only inversion. E,F,G,H: the same, using both stacks. A and B show clear tuning–ambiguity effects in the delineation of layer 3 which are not resolved by the use of the near stack only: an appreciable fraction of realisations permit pinchouts in the thin layer 3. The additional stack resolves this ambiguity (E,F). Insets C and G show shear velocity (v_s) pair-samples of the layer 4 sand and layer 5 shale, showing how the far stack induces the expected shear velocity correlation across layers. The extra constraint does not significantly reduce the uncertainty in most of the histograms of Figure 10, however.

APPENDIX A

CONVERSION OF REGRESSION FORMULAE

Conversion of the regression forms of equations 6 and 7 to that of 8–9 can be derived by simple algebra and assumption of independence of errors. The result is

$$
\begin{aligned}
A_{v_p} &= a_p + b_p(1+g)a_\phi && \text{(A-1)}\\
B_{v_p} &= b_p b_\phi (1+g) && \text{(A-2)}\\
D_{v_p} &= b_p(1 - \frac{1+g}{1-f_c}) && \text{(A-3)}\\
A_\phi &= -\frac{a_p}{b_p(1+g)} && \text{(A-4)}\\
B_\phi &= \frac{1}{b_p(1+g)} && \text{(A-5)}\\
C_\phi &= -\frac{1}{1+g} && \text{(A-6)}\\
\sigma'_{\epsilon_p} &= \sqrt{\sigma^2_{\epsilon_p} + (b_p(1+g))^2 \sigma^2_{\epsilon_\phi}} && \text{(A-7)}\\
\sigma'_{\epsilon_\phi} &= \frac{\sigma_{\epsilon_p}}{|b_p(1+g)|}. && \text{(A-8)}
\end{aligned}
$$

APPENDIX B

DELIVERY IMPLEMENTATION DETAILS

The *Delivery* code works with two versions of the model vector $\mathbf{m}$. The vector $\mathbf{m}$ has a fully Gaussian prior, with no truncations or restrictions on values. The physical model vector $\mathbf{m}'$, which is used in the forward model and its associated likelihoods (seismic, isopachs) is obtained by applying time orderings and truncations (e.g. of NG or saturations) to $\mathbf{m}$, i.e. $\mathbf{m}' = f(\mathbf{m})$, where $f()$ embeds these rules. The truncation effectively induces a prior which, for simple properties like NG, is a mixture of a truncated Gaussian distribution and delta functions at endpoints.

With the augmented models defined by equations 8 and 9, the linearity means the prior is still Gaussian, but the truncation of ϕ_{flt} in $\mathbf{m}'$ must be handled with care. The extra coefficients D_{v_p}, C_ϕ have the effect of placing the prior on inclined ellipsoids in e.g. the $\{v_p, \phi_{\mathrm{flt}}\}$ plane, so pure truncation on ϕ_{flt} has the effect of smearing the tail of the distribution onto the plane $\phi_{\mathrm{flt}} = 0$ in a direction off the principal axes. This is clearly a undesirable way to handle the prior. Figure B-1 shows a scatter plot of points produced from a prior constructed in this naive way, with the obvious artifacts. A more reasonable way to handle the truncation is with the mappings (*only for* $\phi_{\mathrm{flt}} < 0$):

$$
\begin{aligned}
v'_p &= v_p - D_{v_p}\phi_{\mathrm{flt}} && \text{(B-1)}\\
\phi' &= \phi + B_\phi(v'_p - v_p) - C_\phi \phi_{\mathrm{flt}} && \text{(B-2)}\\
v'_s &= v_s + B_{v_s}(v'_p - v_p) && \text{(B-3)}\\
\phi'_{\mathrm{flt}} &= 0, && \text{(B-4)}
\end{aligned}
$$

which forces the remapping to occur along directions parallel to the principal axes.

This mapping minimises the difference $(\mathbf{m'} - \mathbf{m})^T C_P^{-1} (\mathbf{m'} - \mathbf{m})$, subject to the positivity constraint, which seems a reasonable formulation. The prior will then be a mixture of 'clean rocks ($\phi_{\mathrm{flt}} = 0$)' and 'poorly sorted' rocks distributed along the ellipsoid with $\phi_{\mathrm{flt}} > 0$

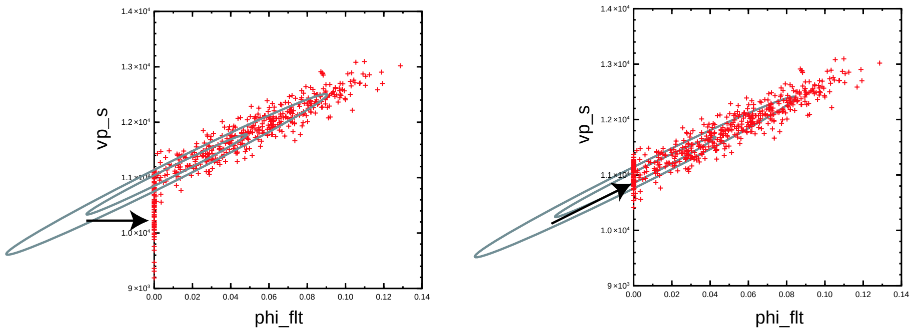

Figure B-1: Left: pure truncation of ϕ_{flt} resulting in smearing of prior along $\phi_{\mathrm{flt}} = 0$ plane. The right figure illustrates the remappings of equations B-1–B-4 which seem more reasonable.

Note that the actual Gassman fluid substitution calculation that occurs later in the forward model uses only the pure fluids (oil, gas etc), as the Gassman–like effect of the floating grain presence is implicitly accounted for by the floating–grain terms in the modified regressions.

REFERENCES

DeMartini, D. C. and M. Glinsky, 2006, A model for variation of velocity versus density trends in porous sedimentary rocks: Journal of Applied Physics, **100**. July, 014910.

Fraley, C. and A. E. Raftery, 2003, Enhanced model-based clustering, density estimation, and discriminant analysis software: Mclust: Journal of Classification, **20**, 263–286.

Gunning, J., 2003, **Delivery** website: follow links from `http://www.petroleum.csiro.au`.

Gunning, J. and M. Glinsky, 2004, Delivery: an open-source model-based Bayesian seismic inversion program: Computers and Geosciences, **30**, 619–636.

Mavko, G., T. Mukerji, and J. Dvorkin, 1998, The rock physics handbook: Cambridge University Press.

Nur, A., D. Marion, and H. Yin, 1991, *in* Hovem, J. M., M. D. Richardson, and R. D. Stoll, eds., Shear waves in marine sediments, 131–140. Kluwer Academic Publishers, Dordrecht, The Netherlands.