Bayesian inversion whispers

Michael E. Glinsky, Mark C. Haase, and Valerie Charoing, BHP Billiton Petroleum, Houston, USA
Guy Duncan, Robin Hill, Gerry O'Halloran, and Long Dang, BHP Billiton Petroleum, Perth, Australia
James Gunning, CSIRO Petroleum, Clayton, Victoria, Australia

Many times we are faced with the business decision of whether or not to develop a sand that is at the limit of seismic resolution and near the noise level of the data. The critical issue is developing a reasonable certainty that there is enough volume of hydrocarbons to develop. A popular approach is to use Bayesian methods to determine the probability of an economic volume of hydrocarbons being present. A problem with this approach when it is applied to these marginal cases is a bias to the answer. Often, this comes from a relatively strong sophomoric prior constraint on the gross thickness and net-to-gross (N/G) of the sands, imposed to keep the inversion focused on the correct seismic reflector. The data are whispering what the answer should be through the Bayesian apparatus, but this whisper is overwhelmed by the sophomoric prior constraints. We found a simple solution to this problem—run the seismic inversion several times using the output mean of the previous inversion as the input mean of the next inversion. This methodology made the difference, in conjunction with a bandwidth improvement in the seismic data, in proving that a well should be drilled. Unfortunately, the well did encounter an acoustically soft lithology of the predicted gross thickness, but it was a shale—the most likely failure mode as predicted predrill.

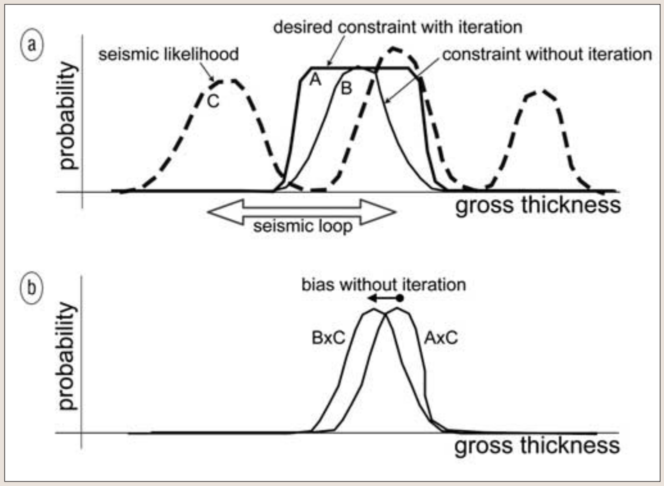

Figure 1. Demonstration of the effect of the prior constraint on the bias of the posterior distribution. (a) A=Prior probability distribution for the desired flat top distribution. B=the practical Gaussian distribution and C=the multimodal seismic likelihood. (b) Posterior probability distribution using the desired flat top distribution, A×C, and the practical Gaussian distribution, B×C.

Introduction to the problem and solution. There has been recent interest and application of Bayesian seismic inversion (Buland and Omre, 2003; Eidsvik et al., 2002; Eidsvik et al., 2004; and Gunning and Glinsky, 2004). The main attractiveness of these methods from a business perspective is the fact that they give an estimate of the uncertainty in the estimate of the volumetrics of potential oil and gas fields. Many times, this is the main contributing factor to economic uncertainty, and therefore can be the deciding factor in business decisions.

A classic decision is whether or not to proceed with construction of an LNG facility. There must be a high degree of confidence that the field will supply enough gas to deliver on the contracts and be economic. This is translated into the uncertainty estimation challenge to have the probability of having the contracted volume to be 90% or greater (i.e., contracted volume less than P90 volume). Often, the outstanding challenge is that the gross thickness of the sands is thin enough so that the magnitude of the reflection is not a lot larger than the seismic noise level. In this situation, Bayesian seismic inversion is helpful in determining the critical uncertainty. That is, advanced technology is used when the answer is not obvious.

The influence of the choice of the prior in these subtle cases can be a problem. The actual value is within a standard deviation or two of the contracted amount so that if the prior is set greater than the actual value by an amount on the order of the posterior standard deviation or more, the posterior estimate of the median and the P90 volume will be biased high by about a standard deviation. This will cause the field to look safely economic when it is not. The opposite is true if the prior is set less than the actual value. There will appear to be significant risk that the field is not economic, when it is probably economic. The decision should not rest on the prior estimate of the volumes.

A purely data-driven way to determine the median and P90, not influenced by the prior assumptions, is needed. The answer is surprisingly simple. Take the mean output of a Bayesian inversion and use that as the mean of the input for a second Bayesian inversion, but do not change the prior standard deviation. Repeat this until there is not much change from the prior to the posterior mean. Through this iterative process, the Bayesian inversion is effectively whispering whether the answer is high or low, allowing the bias to be removed. It is an amazing outcome that the estimate of the mean is biased to much less than a standard deviation.

Another factor that is well known to help resolve the uncertainty of net-sand thickness for thin sands is to increase the bandwidth (effective resolution) of the seismic data. In a sand that is at or below the resolution of the seismic data, the net sand can be determined provided that the sand is acoustically softer than the surrounding shale. Unfortunately, the gross thickness and N/G are not able to be estimated. Once the bandwidth has been increased enough, the gross thickness and N/G can be resolved. This may not solve the problem of the bias if the reflection strength of the resolved sand is not large enough or the noise is too large to see the seismic reflector. The iterative process, in this case, will still be needed to remove the bias in the estimate of the gross thickness, net thickness, and N/G, not only the net thickness of the previous unresolved case.

We show that this iterative process is theoretically convergent and works for a wedge model. We demonstrate the idea on a small target called "Glenridding" under the main pay sand of the Stybarrow Field, offshore Western Australia. We show also the effect of increased resolution of the seis-

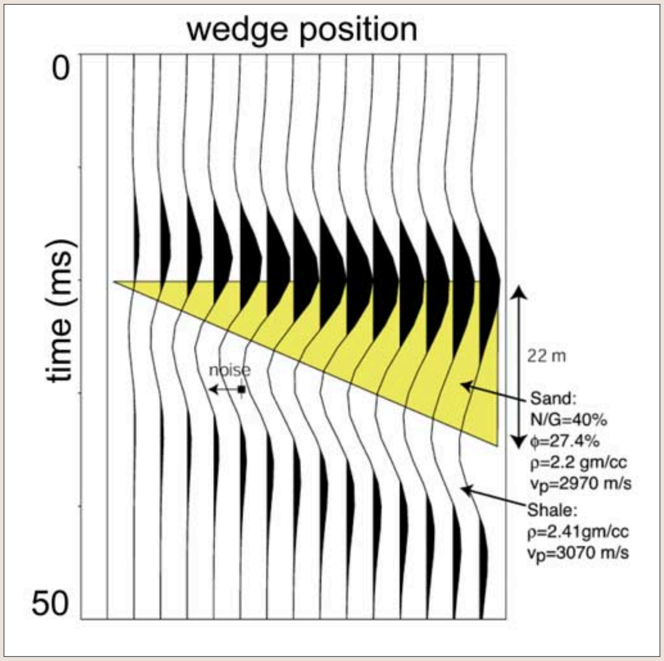

Figure 2. Wedge model and the synthetic seismic. Wiggle traces are reflection coefficient data with a positive amplitude, shaded black, representing a reflection from a hard reflector over a soft reflector.

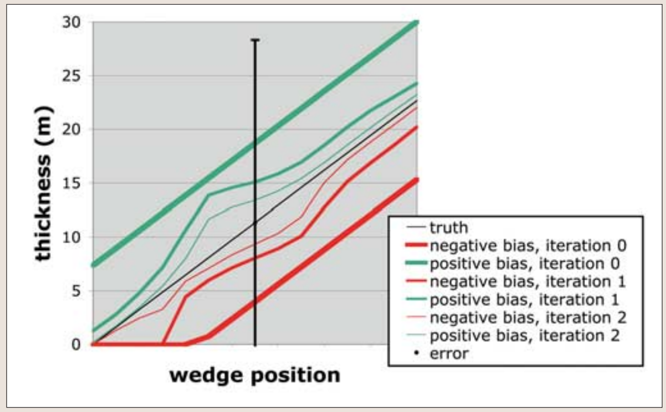

Figure 3. Results of iteration of inversion on the wedge model, starting from two biased cases. Case 1 starts with a positive bias of 5 m (green curves). Case 2 starts with a negative bias of 5 m (red curves). The answer and fixed-point solution is shown as the thin black line. The initial model and two iterations are shown. Both cases converge to the fixed point. The error bar shows a representative standard deviation of the solutions.

mic data for this case. Use of this technology along with improved seismic data will show an accumulation that is probably economic.

Theory of inversion fixed points. We present a general description of the problem and a sketch of the underlying theory, skipping the details of the mathematical proofs. We outline the starting assumptions in a little detail but pass quickly to the practical conclusions of the theory.

This work is based on the Bayesian model-based seismic inversion program (Gunning and Glinsky, 2004). In this layer-based model, a useful prior constraint that makes the inversion problem less multimodal is to focus the inversion on the correct seismic reflection for each sand or shale layer. This is done by imposing Gaussian prior distributions on the layer reflection times, gross thicknesses, and N/G values. These constraints are made as weak as possible, but strong enough to prevent a seismic "loop skip," or trapping in undesirable secondary local minima. This "lion taming" of the objective function is demonstrated by Figure 1. Curve C shows the probability of the model seismic being consistent with the observed seismic as a function of the gross thickness of the sand. The side lobes correspond to the sand becoming thick enough that the reflector would correspond to the top of another sand. The solutions that correspond to these side lobes are not reasonable solutions and need to be excluded. Imposing the Gaussian constraint on gross thickness does this in curve B. Unfortunately, the resulting compound probability, shown as B×C in Figure 1b, has a maximum biased away from the desired local maximum of curve C. What we would like to do is impose a prior constraint similar to curve A with a flat top. This type of nonlinear prior constraint is not used, however, since it breaks the linearity of the inversion, and seriously inflates the numerical demands of the inversion.

To harness the computational advantages of the Gaussian prior (curve B), and the unbiased character of curve A, we implemented an iterative inversion. We start with the normal Gaussian constraint on reflection times, gross thicknesses, and N/G and do the Bayesian inversion. We take the resulting posterior estimates of the mean times, gross thicknesses, and N/G and use them as the mean of the prior constraints (standard deviations remain unchanged) for the next Bayesian inversion. We repeat this process and stop when there is little difference between the prior and posterior means of the three properties.

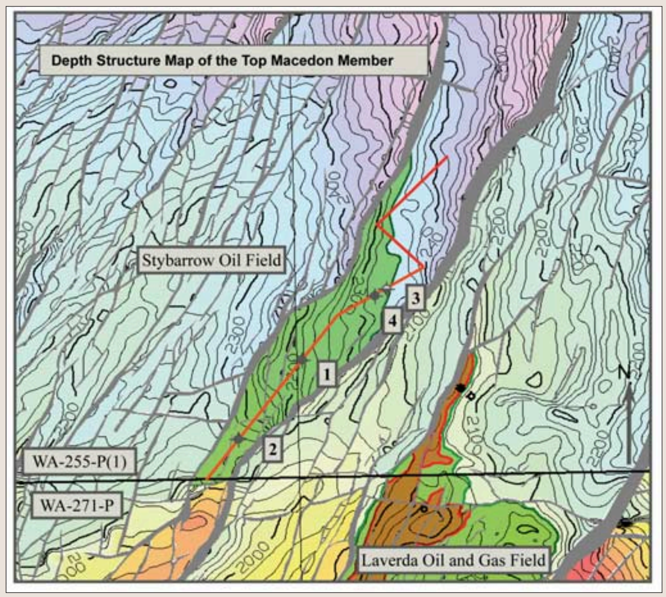

Figure 4. Map of Stybarrow Field with depth contours. The locations of the four appraisal wells are shown. The cross-section location, shown in Figure 5, is indicated by the red line.

In order for this process to work, the mapping of the prior to the posterior means must be a compact mapping whose fixed point has an effective constraint which will not bias the solution, as shown by curve A in Figure 1a. By linearizing the solution about the optimum point (as per Equations 36 and 37 of Gunning and Glinsky), and by separating the prior constraint into the parts that will be iteratively updated and those which will not, it can be proved that a fixed point exists and that the convergence is linear. The fixed point is the solution to the problem with the prior

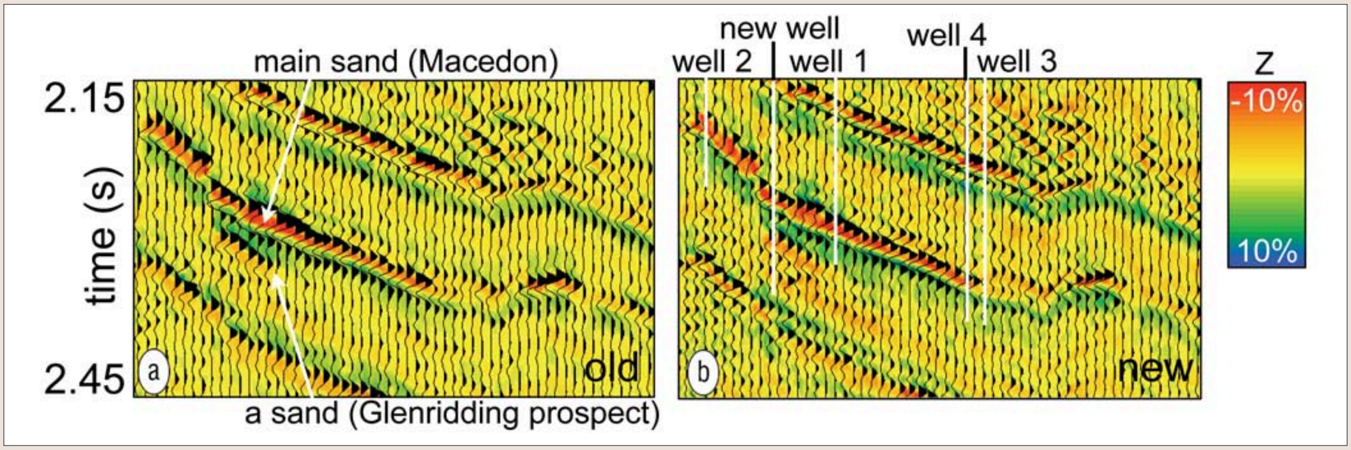

Figure 5. Cross-sections of Stybarrow Field. Seismic reflection data are shown as black wiggles. A positive amplitude is shaded black and represents a reflection from an acoustically hard rock over a soft one. The colored background is a sparse spike inversion. Values are normalized and expressed in relative percent reflectivity. Red is acoustically soft, blue is hard. (a) Old data. (b) New, higher bandwidth data.

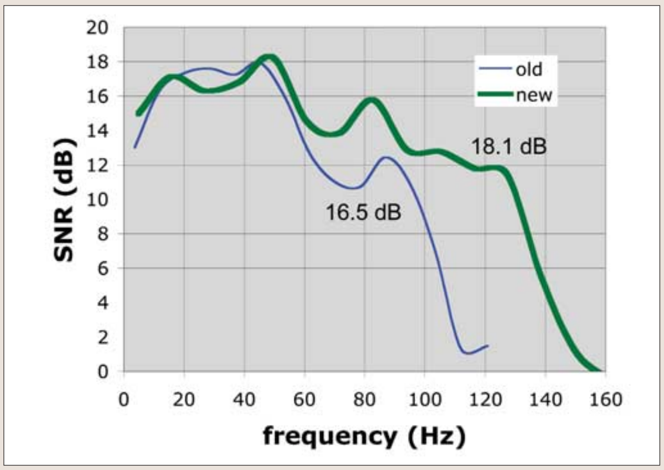

Figure 6. Spectral signal-to-noise ratio determined by comparison to the well log synthetic seismic. Signal is the peak amplitude of the main sand reflector. Posted signal-to-noise ratios are the peak reflection amplitude of the main sand divided by the noise level estimated by the stochastic wavelet derivation (Gunning and Glinsky, 2006).

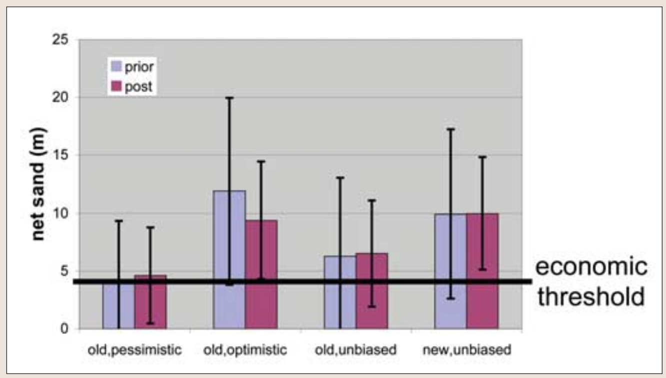

Figure 8. Summary of the estimates of net sand (a sand). Error bars represent the standard deviation. Old, pessimistic case is using the old data with an initial N/G of 85% and an initial gross thickness of 3 m. Old, optimistic case is using the old data with an initial N/G of 85% and an initial gross thickness of 16 m. Old, unbiased case is using the old data with an initial N/G of 85% and an initial gross thickness of 8 m. New, unbiased case is using the new data with an initial N/G of 50% and an initial gross thickness of 22 m.

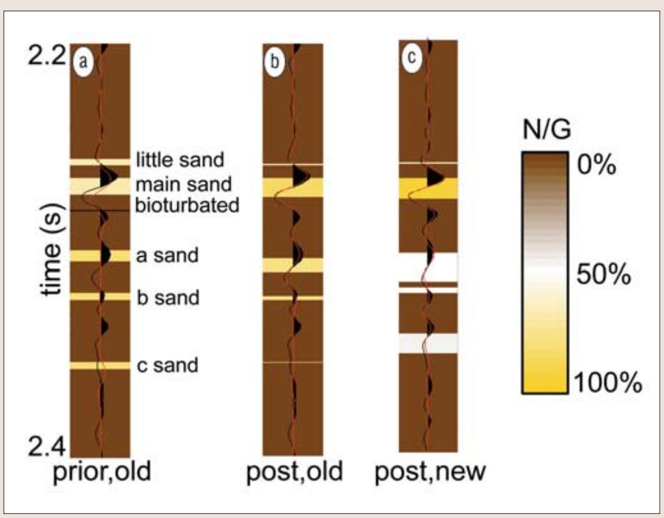

Figure 7. Mean models of N/G. Black wiggle shaded black is the seismic reflection data. Red wiggle is the synthetic seismic of an average model (as determined by the seismic chi squared value). (a) Initial model, before inversion, using the old data. (b) Model after inversion using the old data. Initial N/G was 85%, and initial gross thickness was 8 m. (c) Model after inversion using the new data. Initial N/G was 50%, and initial gross thickness was 22 m.

Gaussian constraints removed on the updated parts. This is exactly what is needed. A condition on the convergence is that the linearized problem is not rank deficient, that is, it is well-posed. This will be true as long as the sensitivity matrix for all the parameters being iterated is "full rank." This means that care should be taken in the choice of the properties that are iterated—the inversion should be significantly decreasing the standard deviations of those properties and they should not be linearly dependent upon each other. Details of the proof can be worked out from well-known results in Tarantola (1987), and Golub and Van Loan (1996).

Simple wedge model. To demonstrate and verify the unbiased result of the iterative inversion, a simple wedge model was constructed. It consists of three layers: a laminated reservoir sand between two shales (Figure 2). The wedge starts at zero gross thickness and linearly increases to a thickness of 22 m. The sand is softer than the shale, and has a N/G of 40%. The end member sand has a porosity of 27.4%, a density of 2.2 gm/cc, and a compressional velocity of 2970 m/s. The shale has a density of 2.41 gm/cc, and a compressional velocity of 3070 m/s. Convolving a Ricker wavelet with the contrast in the acoustic impedance forms the syn-

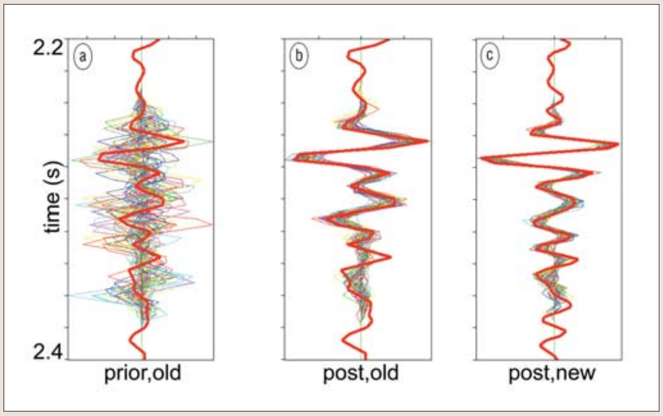

Figure 9. Comparison of synthetic seismic of the ensemble of models to the actual seismic (thick red line). (a) Initial model, before inversion, using the old data. (b) Model after inversion using the old data. (c) Model after inversion using the new data.

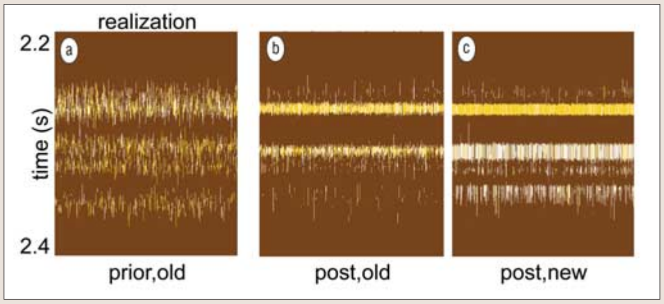

Figure 10. Ensemble of N/G models. (a) Initial model, before inversion, using the old data. (b) Model after inversion using the old data. (c) Model after inversion using the new data. N/G color scale bar is same as that for Figure 7.

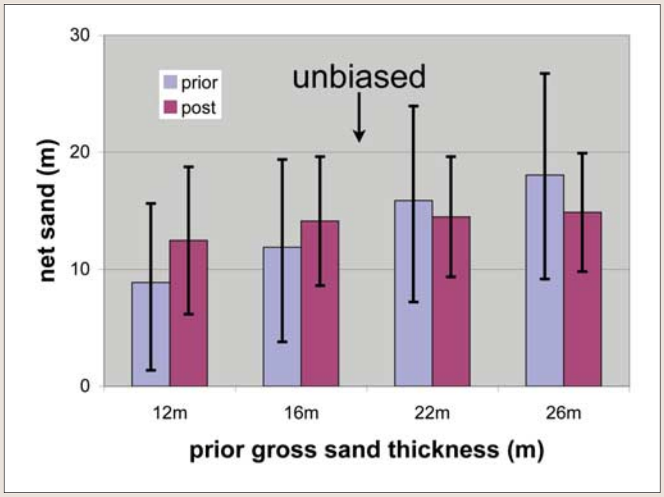

Figure 11. Update to mean net sand (a sand) by the inversion as a function of the initial gross sand thickness. Error bars represent the standard deviation. Initial N/G was 85%. Value indicated as unbiased is the value of the prior gross sand thickness that results in no change of the posterior mean net sand when compared to the prior mean value.

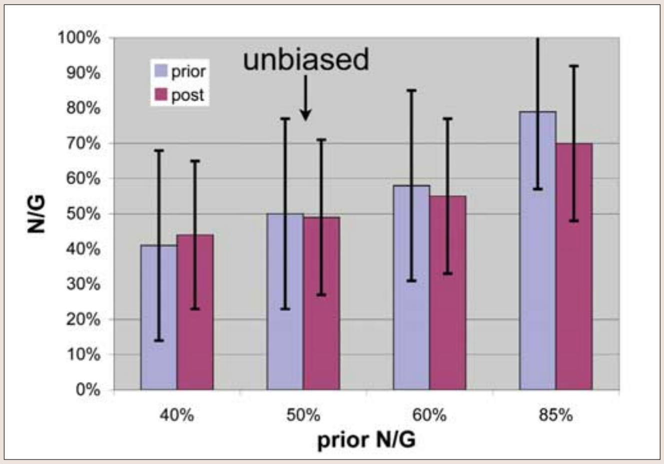

Figure 12. Update to N/G (a sand) by the inversion as a function of the initial N/G. Error bars represent the standard deviation. Initial gross thickness was 22 m. Value indicated as unbiased is the value of the prior N/G that results in no change of the posterior mean N/G when compared to the prior mean value.

thetic seismic. The frequency of the wavelet was chosen to have a tuning thickness of 14 m. The N/G standard deviation for the reservoir layer was 20%. Uncertainties for the end member properties were 87 m/s, 142 m/s, and 1.7% for the sand compressional velocity, shear velocity, and porosity, respectively; and 138 m/s, 70 m/s, and 0.035 gm/cc for the shale compressional velocity, shear velocity, and density, respectively. Time uncertainties of 10 ms were assumed for the two seismic reflectors.

Two series of inversions were done: the first starting with a model that always had 5 m more sand than the model used to construct the seismic, the second starting with a model that always had 5 m less sand. A noise level of about half the size of the reflector was assumed. The gross thickness was iterated for each series of inversions until the solution converged. The result is shown in Figure 3. The posterior uncertainty in the gross thickness was about three times the size of the initial bias. The solutions obviously converged to the unbiased solution. It was noted that the convergence was quite rapid (within one to two iterations) for a noise level merely 50% of that displayed in Figure 2.

Field example. The methodology was then applied to the Glenridding prospect, which lies beneath the Stybarrow Field, influencing the business decision to drill. The Stybarrow Field is in Production License WA-32-L, 135 km west of Onslow, offshore Western Australia. The water depth at the location is approximately 800 m. The field lies near the southern margin of the Exmouth sub-basin within the larger Carnarvon Basin. Oil is trapped in the Early Cretaceous, Berriasian age turbidite and debris-flow sandstones deposited on a relatively shallow passive margin slope. The Stybarrow structure comprises a NE–SW tilted fault block, forming a terrace within the westward plunging Ningaloo Arch (Figure 4). The intersection of SW–NE and E–W normal faults establishes an elongate, triangular trap forming structural closure to the southwest. The structure dips from the SW to the NE at about 5°. Top, base, and bounding fault seals are provided by claystones and siltstones of the overlying Muiron member of the Barrow group and mudstones of the underlying Dupuy Formation. More information about the field can be found in Ementon et al. (2004).

A seismic dip cross-section through the middle of this field is shown in Figure 5. Note the main sand that is currently under development (the Macedon sandstone) and the location of the four appraisal wells. The seismic data were recently reprocessed in a way that increased the bandwidth (Figure 6). This reprocessing highlighted a small, but possibly economic "a sand" called the Glenridding prospect, that has not yet been penetrated, approximately 50 m below the main Macedon sand

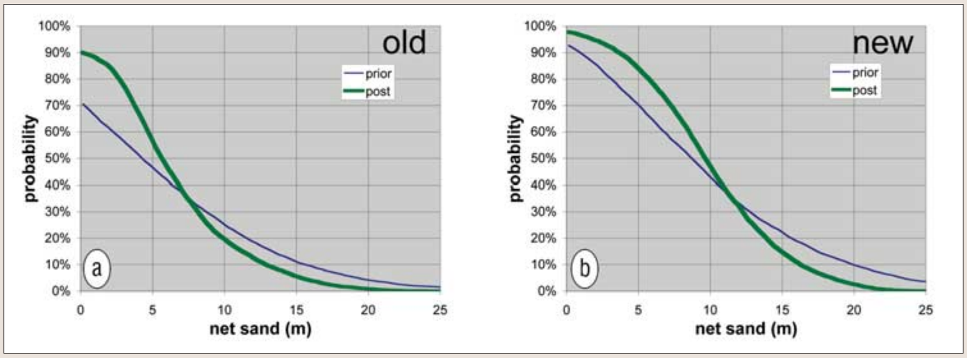

Figure 13. *Cumulative probability distributions of net sand. (a) Old seismic data. (b) New seismic data.*

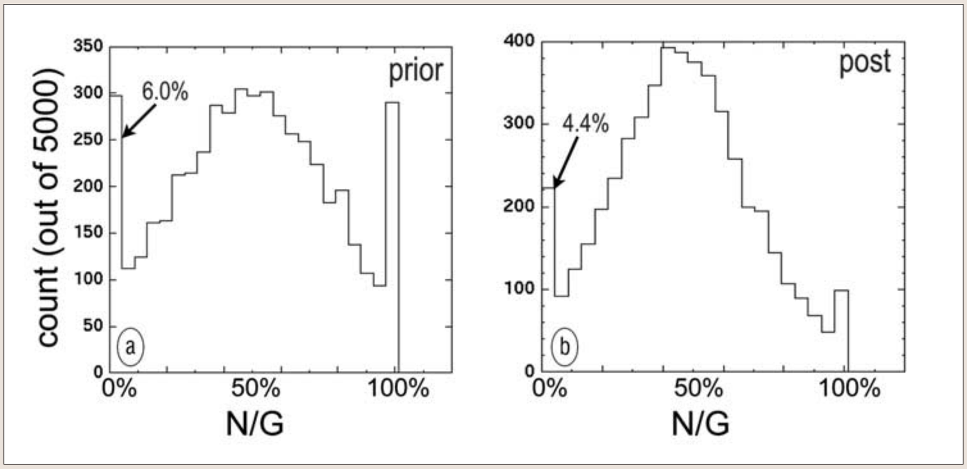

Figure 14. *Probability distributions for N/G for the fixed point using the new seismic data. (a) Before inversion. (b) After inversion.*

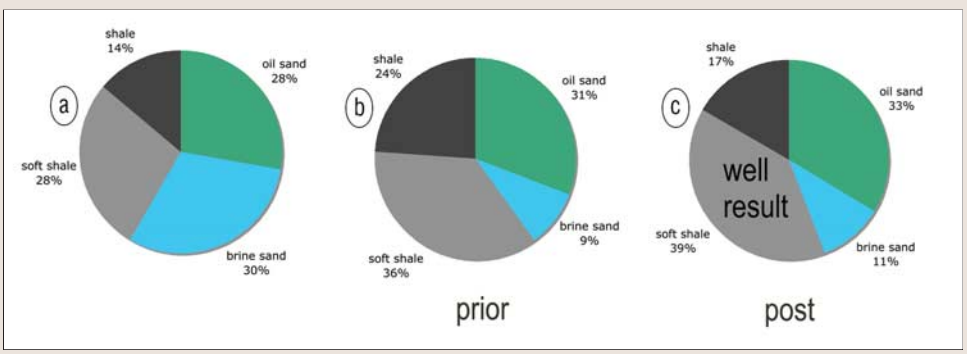

Figure 15. *Wheels of fortune for the four possible models. (a) As indicated by the seismic data assuming equal prior probabilities for the four models. (b) Prior probabilities as determined by the asset team considering the well information, geology and petroleum system. (c) As indicated by the seismic data assuming the asset team's prior probabilities.*

(see Figure 5). Given the limited areal extent of this near-field prospect, it would need to have a thickness of at least 4 m to break even economically. In order to drill this target, it needs to be proven that there is a 90% probability of having at least 4 m of sand.

To answer this question, a Bayesian model-based inversion was done at the proposed well location shown in Figure 5. A model was constructed as shown in Figure 7. It has 12 layers, five of which are sands. It was built from an interpretation of the top and base of the main (Macedon) sand, and the top of the Glenridding "a sand." Small uncertainty was assumed for the position of these interpreted horizons (6 ms), and a larger uncertainty for the other horizons (8 ms). The uncertainty in the N/G was assumed to be 30% with a initial mean of 85% for all sands. More details on how this inversion was done can be found in Glinsky et al. (2005).

The first inversion was done using the old data. The seismic specialists doing the inversion decided to make a pessimistic assumption for the initial thickness of the "a sand"—they assumed that it had zero thickness and had the inversion prove otherwise. Because the noise level was about the size of the seismic reflection, this was a reasonable possibility. The resulting estimate of the net sand (Figure 8) does not meet the criteria for drilling the well. The asset team members challenged this result, suggesting an optimistic assumption that there is a sand of tuning thickness unless proven otherwise. This was also a reasonable possibility. The resulting estimate, also shown in Figure 8, does meet the criteria for drilling the well. Who was right? Finding the answer to this question was the inspiration for the discovery of the iterative inversion. The unbiased answer using the old and the new data is shown in Figure 8. The unbiased solution using the old data is obviously a compromise between the optimistic and pessimistic solutions and unfortunately does not meet the criteria for drilling the well. Fortunately, the resolution provided by the new data increases the estimate of the net sand enough to meet the criteria for drilling the well. Note that the new, unbiased result is consistent with the old, unbiased result (i.e., the new mean lies within the uncertainty of the old data), but it is not consistent with the old, pessimistic result.

Let us now examine the results in more detail so that we can better understand them. Start with the mean models shown in Figure 7. Both inversions using the old and the new data increase the N/G and gross thickness of the main sand. There is no change to the N/G of the "a sand" using the old data and a very modest increase to the net sand. This is because this sand is not resolved. The new data set is able to resolve this sand. It dramatically increases the thickness, but significantly decreases the N/G with an increase in the net sand. It also increases the N/G of the main sand. The match of the model synthetic seismic to the seismic is shown in Figure 9. Note the better match using the new data due to the better signal-to-noise ratio (SNR). A very instructive perspective on the inversion results is obtained by examining all of the possible models that fit the data to within the SNR (Figure 10). Notice the reduction in the uncertainty in the location of the top and base of the main and "a" sands with the new data. There is also less scatter in the N/G of both sands using the new data.

The existence of the fixed point and the convergence can be seen in Figures 11 and 12. They show the change in the net sand and N/G, respectively, as the prior mean values are changed. Note that, for values less than the fixed point (labeled unbiased), the inversion increases the value. The greater the distance from the fixed point the larger the change. The opposite is true for values greater than the fixed point—the inversion decreases the value.

The bottom line results are shown in Figure 13, where the cumulative distribution functions are shown for the net sand in the "a sand" using the old and the new data. Note that there is only a 65% probability of having at least 4 m of sand using the old data, and that probability is increased to 90% using the new data.

Epilog to field example. The Glenridding well has recently been drilled, subsequent to this analysis. Unfortunately, it found an acoustically soft shale more than 13-m thick (the predicted gross thickness by the inversion was 19±7 m). The compressional velocity of this shale was 300 m/s slower (two standard deviations) than expected for a shale at this depth. Encountering a soft shale was a concern before the well was drilled since a soft shale was penetrated in an equivalent stratigraphic interval in a well 30 km away from Stybarrow.

This is a surprise given the previously presented analysis of this paper. A clue to what the problem is can be seen in Figure 14. It shows the N/G distributions for the target sand before and after the seismic inversion. Note that there was only a 6% chance of no sand being present in the prior distribution. Although the probability was reduced 25% to 4.4% by the seismic inversion, the seismic data were not definitively eliminating this as a possibility. The small posterior probability was mostly due to prior assumption on N/G.

The solution to this problem is to explicitly consider a brine sand, regular shale, and soft shale as alternative models to the oil sand. This was done by the asset team before the well was drilled. The results are shown in Figure 15. Assuming equal prior probabilities for the four cases, Figure 15a shows that the seismic response is equally consistent with the brine sand, oil sand, and soft shale models. It is less consistent with a regular shale. The prior probabilities of the asset team (considering the nearby well control, geology, and petroleum system) are shown in Figure 15b. When these probabilities are updated with the Bayesian boost from the observed seismic response, the result is Figure 15c. Note that the most likely model is the soft shale (39%) and it is 56% likely that some type of shale would be found. The well result (finding no sand) was therefore a very likely occurrence.

The model selection was done by an extension of the Bayesian model-based inversion program that outputs the marginal model likelihood for the four inversions—one for each model. More information on Bayesian model selection can be found in Gilks et al. (1996) and Denison et al. (2002).

Conclusions. The iterative inversion is an important refinement to Bayesian inversion when looking at marginal sands where seismic reflection amplitude is near the noise level. For high noise levels, the data are whispering to you through the Bayesian inversion, but are being overwhelmed by the heavy-handed sophomoric prior constraints imposed to eliminate unreasonable models. The iteration amplifies the whisper allowing convergence to the unbiased, predictive result free of the influence of the initial constraints. Increasing the bandwidth of the seismic data also is important if these sands are poorly resolved. These methodologies changed a business decision, but unfortunately the most likely failure case was found when the well was drilled.

Suggested reading. "Bayesian linearized AVO inversion" by Buland and Omre (GEOPHYSICS, 2003). "Seismic reservoir prediction using Bayesian integration of rock physics and Markov random fields: A North Sea example" by Eidsvik et al. (*TLE*, 2002). "Stochastic reservoir characterization using prestack seismic data" by Eidsvik et al. (GEOPHYSICS, 2004). "Delivery: An open source model-based Bayesian seismic inversion program" by Gunning and Glinsky (*Computers and Geosciences*, 2004). "Stybarrow oil field—From seismic to production, the integrated story so far" by Ementon et al. (SPE Asia Pacific Oil and Gas Conference, Expanded Abstracts, 2004). "Integration of uncertain subsurface information into multiple reservoir simulation models" by Glinsky et al. (*TLE*, 2005). "Delivery-extractor: A new open source wavelet extraction and well tie program" by Gunning and Glinsky (*Computers and Geosciences*, 2006). *Inverse Problem Theory, Methods for Data Fitting and Model Parameter Estimation* by Tarantola (Elsevier, 1987). *Matrix Computations* by Golub and Van Loan (John Hopkins University Press, 1996). *Markov Chain Monte Carlo in Practice* by Gilks et al. (Chapman and Hall, 1996). *Bayesian Methods for Nonlinear Classification and Regression* by Denison et al. (Wiley, 2002). TLE

Acknowledgments: The authors acknowledge the financial support of the BHP Billiton Technology Program, and thank BHP Billiton and Woodside Petroleum for permission to publish the results.

Corresponding author: mglinsky@bhpbilliton.com