The value of using relative amplitude changes

Michael E. Glinsky, Robert Pascoe, and Bruce Asher, BHP Billiton Petroleum, Houston, USA
Guy Duncan, BHP Billiton Petroleum, Perth, Australia
James Gunning, CSIRO Petroleum, Clayton, Victoria, Australia

There is a rich history of using differential measurements to improve the signal-to-noise (S/N) ratio when there is correlated signal to improve the signal-to-noise ratio. With respect to seismic amplitude measurements, this is done in a practical sense by comparing the amplitude of reflections on structure to those off structure, or by comparing anomalous amplitudes to an average background. In fact, anomalous is defined in reference to the background. While this is intuitively the thing to do, the value of this methodology needs to be calculated by the quantitative effect on risk and uncertainty and the interaction of this effect with business decisions. It is only by influencing the business decisions with the information that value is realized.

We will explore this topic by first giving an introduction to differential measurements and giving a familiar example (telling day from night by the temperature). The hypothesis of how this can be extended into quantitative seismic amplitude analysis will be presented. A successful, calibrated Bayesian framework for how observed amplitude response can be transformed into risk will be explained. This is a prerequisite for calculating the value of using relative amplitude changes. A further Bayesian methodology that incorporates the correlation into a model-based inversion will be used to quantify the estimation of the volumetric uncertainty. Finally the value of this technology will be estimated on a typical exploration example in an undrilled area. A significantly larger value will be shown in a competitive bidding situation against a party without this superior methodology.

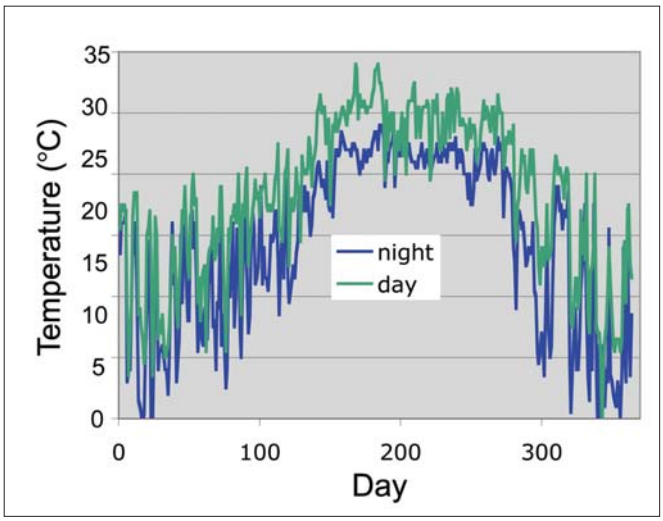

Figure 1. *Plot of the daytime and nighttime temperatures for Livermore, California, during the year 2005. The daytime temperature is taken at 4 P.M. local time. The nighttime temperature at 4 A.M. local time (data courtesy of NOAA Satellite and Information Service, National Climatic Data Center).*

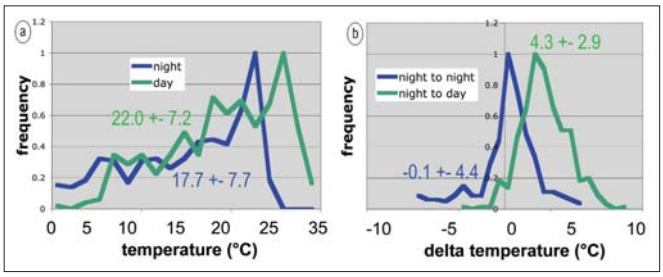

Figure 2. *(a) Distribution of daytime and nighttime temperatures. The mean and standard deviation are shown for each distribution. (b) Distribution of the difference in temperature going from night to the next day, and from one night to the next night.*

Concept of differential measurements. Differential measurements have been used in many disciplines to improve the S/N ratio when a correlation in the signal exists. Recent publicity of the positive drug test of Floyd Landis in the Tour de France highlighted a use of differential measurements in the medical field. Although the amount of testosterone can vary widely in a human being, the ratio of the amount of testosterone to epitestoterone (T/E) is very highly correlated. Therefore, an elevated T/E ratio is used as evidence of an artificial enhancement of testosterone to improve one's athletic performance. In a Bayesian context there is a greater separation in the expected range of T/E given that there has or has not been artificial enhancement, than there is in the amount of testosterone. This allows one to have a greater probability of a person artificially enhancing his testosterone level given an elevated T/E ratio, in comparison to an elevated testosterone level.

There are also several very useful applications of differential measurements in the electronics field. They range from low-voltage differential signaling (i.e., gigabit Ethernet), to USB communication protocols, to high-voltage differential signaling in SCSI-1 equipment.

An Earth science example that we will explore in some detail is the ability to tell day from night using temperature. Figure 1 shows the daytime and nighttime temperatures for Livermore, California, every 24 hours for the year 2005. When the distribution of daytime and nighttime temperatures is examined in Figure 2a, there is a very large overlap. The standard deviation is about 7°C while the separation in the mean of the two distributions is only 4.5°C. By examining the relative change in temperature over a 12-hour period (daytime) and a 24-hour period (nighttime), there is a much larger relative separation in the distributions (Figure 2b). While the separation in the mean of the distributions remains 4.5°C, the width of the distribution is decreased by almost a factor of 2 to 4°C. This has a profound effect on the ability to tell daytime from nighttime. Using a Bayesian analysis where the probability of it being daytime is 50% and assuming that there is an observation of either the most likely daytime temperature or 12-hour change in temperature, Figure 3 shows the advantage of using the relative temperature change. Given only the observation of the temperature, the proba-

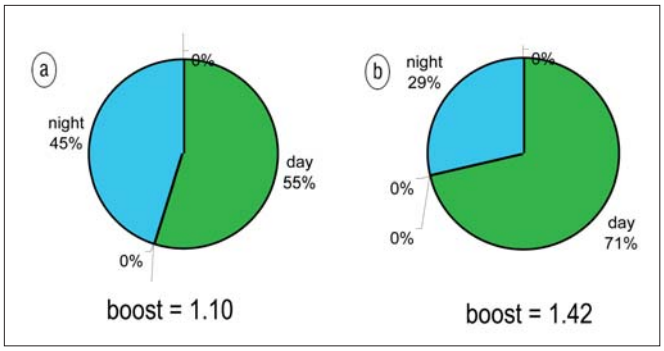

Figure 3. *(a) The probability of day and night given the observation of the mean daytime temperature. The prior probabilities are assumed to be equal for day and night. The boost is the probability of day given the observation divided by the prior probability of day. (b) The probability of day and night given the observation of the mean difference between night and day.*

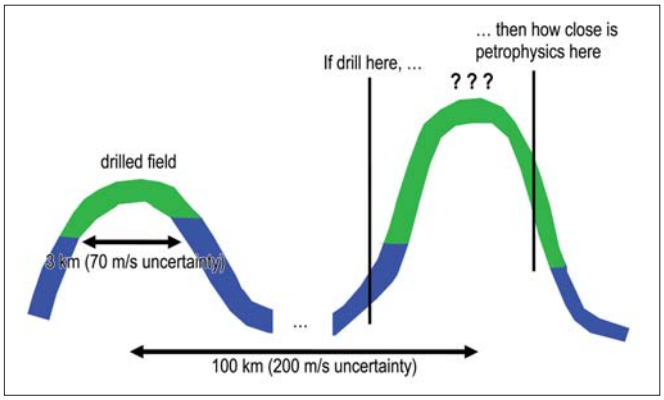

Figure 4. *Context of petrophysical correlation. The small uncertainty over a drilled field is shown on the left. There is a larger uncertainty going to an undrilled field many kilometers away. Although the petrophysical properties at that location are quite uncertain, the properties downdip should be similar to the updip properties.*

bility that it is daytime is only 55%, a boost in the probability of 10%. The observation of the temperature change increases the probability that it is daytime to 71%, a boost in the probability of 42%. Obviously the use of the differential temperature measurement increases the predictability because of the short time (daily) correlation in temperature.

The differential measurement that is the subject of this paper is the local change in seismic reflection amplitude between the updip portion of a potential hydrocarbon trap and the downdip portion of the same trap (Figure 4). In an exploration situation where one only has petrophysical information on the rock physics from wells 100 km away from the potential hydocarbon trap, it is not unusual to have an uncertainty in the compressional velocity of the end member sands and shales of 200 m/s. If one were to drill a downdip well in the potential hydrocarbon trap and derive the compressional velocity of the end members from that well, it would not be uncommon to have an uncertainty of only 70 m/s in predicting the updip end member velocities (over a distance of less than 3 km). The larger variation of the regional prediction is analogous to the seasonal fluctuation of temperatures in the previous example, while the smaller variation within the potential hydrocarbon trap is analogous to the daily variation in the temperature. The local correlation in the petrophysical properties should allow a better prediction of the presence of hydrocarbons using the local change in seismic reflection amplitude rather than using the absolute seismic reflection amplitude. The remainder of this paper will quantify the improvement in the prediction, and its business value.

Bayesian risk determination. This information must be put in the context of an integrated and calibrated risking system to quantify the effect of the relative amplitudes on the risk and thereby the value of hydrocarbon prospects. The methodology that we choose to use is Bayesian-based.

It assumes that there are several distinct classes that need to be distinguished. Within the context of this paper, it is assumed that there are only three classes: shale only (no reservoir), reservoir with brine sand surrounded by shale, and reservoir with hydrocarbon (oil). The prior probabilities of these classes are determined by standard geologic risking based on reservoir presence, existence of a trapping configuration, sufficient sealing capability of the trap, and the likelihood of charge (source, migration, and timing). Additional factors such as the probability of biodegradation and the likelihood of oil versus gas are considered, as needed.

Utilizing knowledge of the petrophysical properties of the end member sands and shales (density, compressional velocity, and shear velocity) along with the properties of the fluids and the N/G of the sands, a range of possible amplitude responses are generated for each class. The uncertainty of all the input properties obviously needs to be known in order to do this. Given the observed amplitude response of the seismic data, the probability of hydrocarbon along with the most likely volumetrics can be estimated. More details can be found in both Duncan et al. (2004) and Glinsky et al. (2005).

Two additional factors need to be taken into account before the risk can be determined. The first is the confidence in the seismic and petrophysical data. One should ignore the seismic information if these data are fundamentally flawed so that the amplitudes cannot be trusted. We estimate the probability of this happening by a deterministic formula based on answers to qualitative questions relating to data quality. This confidence is then used to discount the Bayesian update to the probabilities due to the observation of the amplitudes. The final factor is the fit to structure of the amplitudes; that is, whether gravity is at work on the segregation of the brine and hydrocarbons. This is incorporated as an additional Bayesian update based on the quality of the structural fit. We refer you to Glinsky et al., 2004 for more details on this process. The final result is shown as the update to the "wheel of fortune" similar to Figure 3.

The most important part of this process is the calibration of the results. Over the period of 2000–2005, this methodology was applied before 22 wells were drilled. The expected number of successes ranged from 9 to 16 with a most likely value of 12.5. There were 12 successes, indicating that the method is well calibrated (Figure 5).

Addition of correlation to modeling. In order to study the effect of the correlation in the seismic reflection amplitude on the risk and uncertainty, the Bayesian model-based inversion program DELIVERY (Gunning and Glinsky, 2004) was modified to include this correlation. Instead of correlating the properties at different locations, the problem was simplified by correlating the properties of different layers in the model. Figure 6 shows the net-to-gross-sand property for the prototype model that will be examined extensively in this paper. The upper sand layer is a surrogate for the updip location, and the lower sand layer is a surrogate for the downdip sand location. The middle shale layer is actually made up of three layers. This is done to allow the shale on top of the downdip sand to be different than the shale below the updip

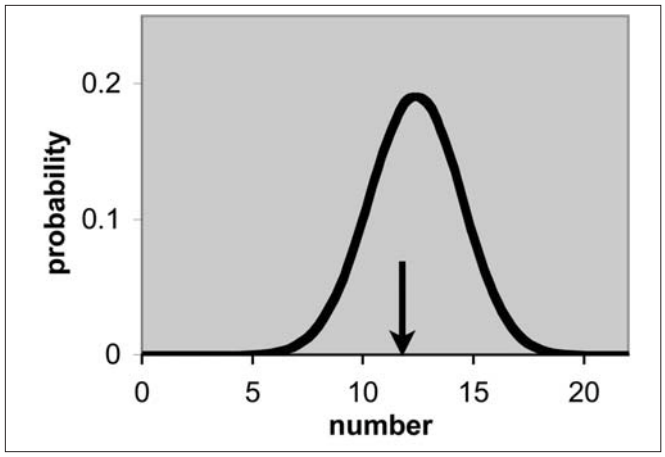

Figure 5. Calibration of risking. Black curve is the Poisson distribution of the expected number of wells that should encounter hydrocarbons. This was for the 22 wells drilled by BHP Billiton between 2000 and 2005 that had an amplitude response that could be analyzed. The black arrow shows the actual number of wells which encountered hydrocarbons, 12.

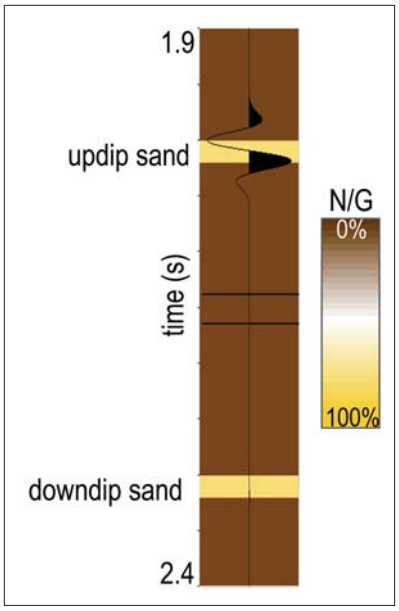

Figure 6. The seven-layer model used in the analysis. Displayed is the N/G of those layers, colored according to N/G. The top three layers are correlated to the corresponding bottom three layers. The upper sand is a surrogate for the sand at the updip location. The lower sand is a surrogate for the sand at the downdip location.

sand. The reflection of the shale between these two shales is ignored by the inversion.

As is customary in DELIVERY, the properties of the end members are specified with uncertainty. What has been added for this study is a correlation between equivalent layers, for instance the updip to the downdip sand and the shale on top of the updip sand to the shale on top of the downdip sand. The specific relationships that were used for this study are for sands:

$$V_P = 1564 + 0.9804\,D \pm (210, 75\ \text{m/s}) \quad (1)$$
$$\phi = 0.6530 - 1.33 \times 10^{-4}\,V_P \pm (0.05, 0.015) \quad (2)$$
$$V_S = -1016 + 0.8454\,V_P \pm (122, 61\ \text{m/s}) \quad (3)$$

and for shales:

$$V_P = 1329 + 0.9448\,D \pm (192, 76\ \text{m/s}) \quad (4)$$
$$\rho = 0.2038\,V_P^{0.3120} \pm (0.04/0.015) \quad (5)$$
$$V_S = -627 + 0.7027\,V_P \pm (122, 61\ \text{m/s}) \quad (6)$$

where V_P is the compressional velocity (in m/s), V_S is the shear velocity (in m/s), ϕ is the porosity, ρ is the density (in g/cc), and D is the depth below the seafloor (in m). The numbers in parentheses are the uncertainties (standard deviations). The first number is the regional uncertainty, while the second number is the local uncertainty within the potential hydrocarbon trap. The depth (D) assumed for the study is 2000 m. We have also introduced a correlation in the gross thickness of the sand, 30 ± (18 m, 5 m), and in the ratio of the net-to-gross thickness, N/G = 60% ± (30%, 10%). The underlying model that we use to generate the synthetic seismic response will have an N/G of 80% and a gross thickness of 32.6 m. All of the end member rock properties will be their most likely values, which give a porosity for the sand of 18.2%. The uncertainty in N/G around the most likely value of 80% will be (15%, 7.5%). The prior distributions for both the N/G and gross thickness input to the model-based inversion has been chosen to cover both the actual values and a value of 0. The fluid properties are assumed to be 1640 m/s and 1.00 g/cc for the brine, and 940 m/s and 0.63 g/cc for the 100% oil. The water saturation for the oil is assumed to be 30%.

The reduced standard deviation of the difference of two properties that are physically proximate occurs because of

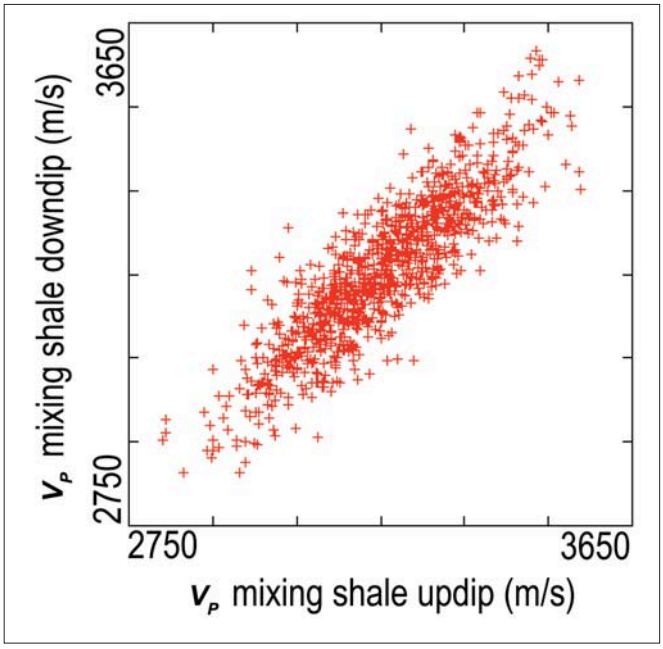

Figure 7. Demonstration of the correlation in the rock properties. Shown is the compressional velocity of the mixing shale used in the two sand layers. While there is a large scatter along either axis, there is a much smaller scatter along the direction indicating the difference between the two properties.

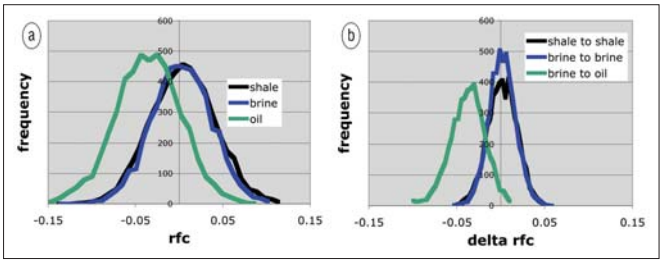

Figure 8. (a) Distribution of reflection coefficients of the top of the upper sand given that the fluid in the pore space is brine or oil. The third case is if there is no sand in the potential reservoir interval. The reflection in this case is caused by a different type of shale. (b) Distribution of the change in reflection coefficients going from the top of the downdip sand to the top of the updip sand.

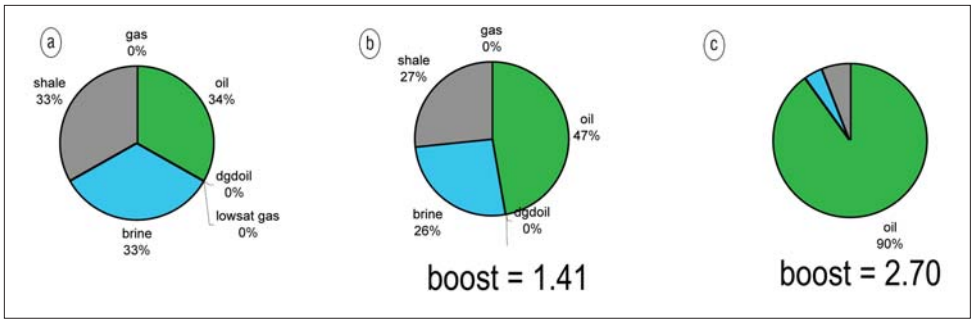

Figure 9. *The effect of observing the mean oil response on the probabilities of each of the three cases: (a) equal prior probabilities before consideration of the seismic amplitude response, (b) the result of consideration of the absolute amplitude of the upper sand, (c) the result of consideration of relative downdip to updip change in amplitude response. The boost is the ratio of the oil probability after consideration of the amplitude response to the prior probability of oil.*

two effects. The first is that the "univariate" standard deviation obtained from a regional or basin-wide trend will be derived from data that are often a mixture of subpopulations, e.g., mixed depositional environments, varied grain characteristics, variations in source sediment etc., which make the overall population variance large. The updip-downdip comparison will be applied at two locations sufficiently close to be described by only one of the subpopulations that comprise the regional trend, and the associated within-group variance is much reduced. Secondly, any "random" spatial process for one of the subpopulations will have some characteristic transverse correlation length, and the difference of values measured well within this correlation length has a standard deviation less than that of the subgroup (Deutsch, 2002). The first effect in particular can be dramatic and can lead to factors of 2 or 3 in the ratio of standard deviations. There have been some interesting studies of the second effect by Eidsvik et al. (2004) and Buland et al. (2003).

The effect of the correlation can be seen in Figure 7 which shows the compressional velocity of the updip mixing shale crossplotted against the compressional velocity of the downdip mixing shale. Note that the spread of the values along either axis is about 200 m/s while the spread along the axis diagonally from the upper left hand corner to the lower right hand corner (the direction that corresponds to the difference between the two velocities) is less than half. This is the expression of the correlation of the updip to downdip rock properties.

Results of the modeling with correlation. We will study two cases to understand the effect of the rock physics correlation and differential measurements on the risk and volumetric uncertainty estimation. The first is the basic situation of the reflection off of the interface between the two infinite half spaces, consisting of a shale overlying a laminated sand-shale mixture. This will be modeled by examining the reflection coefficients of the model shown in Figure 6. Three cases are analyzed. The first is where the sand has 0% N/G (shale). The second is where there is a brine sand both updip and downdip. The third is the same as the second model except with oil updip and brine downdip.

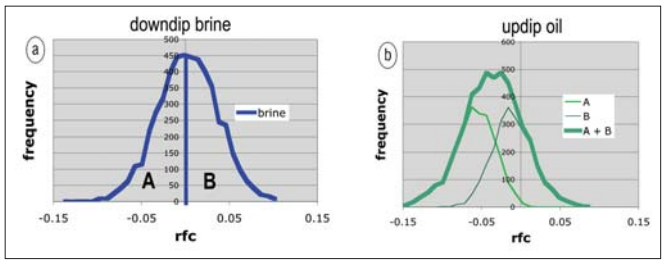

Figure 10. *Demonstration of effect of correlation on segregation of updip oil reflection strength. (a) Division of downdip brine sand amplitude response into two parts: A, soft brine sands, and, B, hard brine sands. (b) The segregation of the reflection coefficients of the correlated updip oil sand. The composite distribution is the same as shown in Figure 8a but is made up of two distributions—one associated with the soft downdip brine sands that is acoustically very soft, and another associated with the hard downdip brine sands that is nearly transparent on average.*

The distributions of the reflection coefficients of the updip "sands" are shown in Figure 8a. In contrast, the distributions of the differential reflection coefficients (top of the updip "sands" less the top of the downdip "sands") are shown in Figure 8b. The feature to note is the better separation of the peaks using the differential reflection coefficients, analogous to what we found with the temperature example. The consequence of this separation is a better ability to identify the presence of hydrocarbons present, in other words, a risk reduction. This is demonstrated by Figure 9 which shows the effect of observing the most likely seismic amplitude response for an oil sand (third situation modeled) on the risk, starting

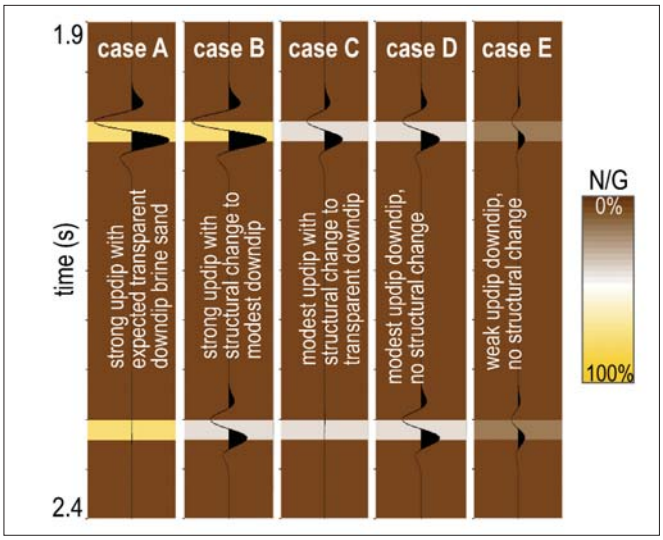

Figure 11. *Characteristic downdip to updip seismic responses whose implication on risk and volumetric uncertainty is analyzed via a Bayesian model-based inversion. The seismic trace is shown as the wiggle with the positive values colored black. They are zero phase with a negative indicating a reflection from a transition into an acoustically softer rock. The N/G of the model is shown as the color behind the seismic wiggles. Shown are five cases: (a) strong updip reflector with the expected structural change to a transparent downdip brine sand reflector, (b) strong updip reflector with a possible change to a modest soft downdip brine sand reflector, (c) modest updip reflector with possible change to a transparent downdip brine sand reflector, (d) modest updip and downdip reflector with no structural change, (e) weak updip and downdip reflectors with no structural change.*

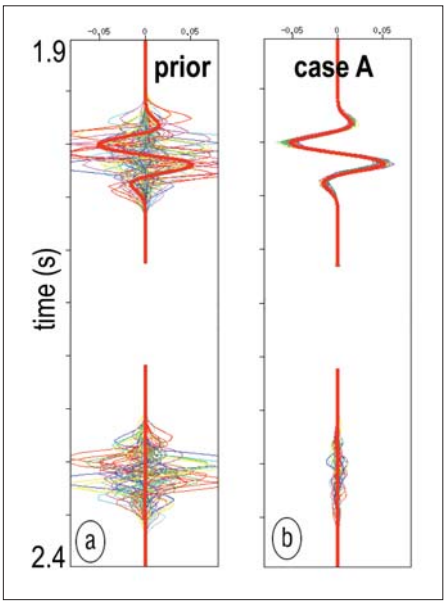

Figure 12. *The result of the Bayesian model-based inversion on the ensemble of synthetic seismic of the models. The "observed seismic" is shown as the thick red trace. It corresponds to case A. The thin multicolored traces are the synthetic seismic of the ensemble of models. (a) Before taking the observed seismic into account. (b) After making the models consistent with the observed seismic to within the noise level.*

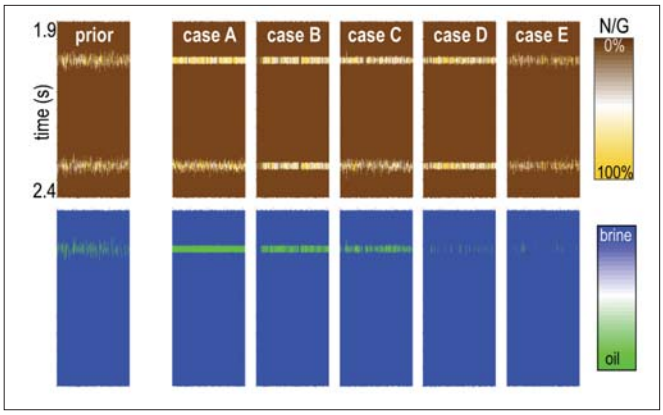

Figure 13. *Ensemble of models. On the top are the N/G of the models and on the bottom are the fluids in the sand. Shown is the ensemble prior to making it consistent with the observed seismic, next are the ensembles made consistent with each of the observed seismic cases shown in Figure 11.*

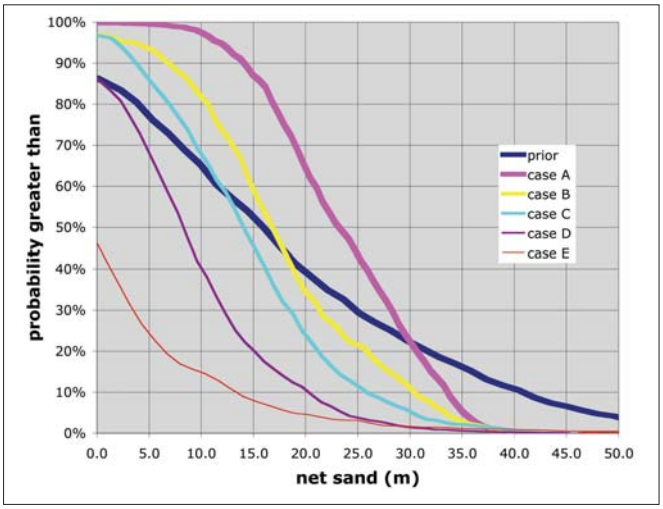

Figure 14. *The cumulative distribution function of the net sand (N/G times gross sand thickness) for each of the cases shown in Figure 11. Also shown is the cumulative distribution prior to making it consistent with the observed seismic.*

from equal probabilities for the three situations (shale, brine, and oil). The probability of oil is only increased from 34% to 47% by the absolute updip amplitude—a relative increase in probability of 41%. This same probability is increased to 90% from 34% using the differential seismic amplitude change—a relative increase in probability of 170%.

To better understand what is causing the tighter distribution in the differential amplitudes and hence the reduction in risk, examine Figure 10a. It divides the brine peak of Figure 8a (downdip sand) into two halves—a sand that is softer than the bounding shale and a shale that is harder than the bounding shale. The reflection coefficients of the updip sand are then calculated corresponding to the two populations of the downdip sand. Because of the rock property correlation, the soft downdip brine sand corresponds to the very soft updip oil sand because of the rock property correlation (Figure 10b). The hard downdip brine sand corresponds to a mildly soft updip oil sand. When the two distributions for the updip oil sand are added together, one gets the broader distribution corresponding to the distribution of the absolute amplitude of the oil sand shown in Figure 8a.

The second case that we will study is the more general case of a three-layer model that has seismic interference between the top and the base sand reflector. This will allow us to study the effect of the tuning and the ability to reduce the uncertainty in both the net and gross thicknesses. The importance of estimating the range of net sand uncertainty comes from its direct relationship to the volume of hydrocarbon. There are five prototype situations that we will examine (Figure 11). They range from the expected response for an oil sand updip and a brine sand downdip to the expected response for shale. The most common cases that we tend to risk are case B (a reasonable amplitude change with a mappable downdip sand) and case D (a mappable updip and downdip sand with little change in amplitude as one moves downdip). A common excuse for the lack of amplitudes when case D is observed is that the low porosity of the sands (18%)

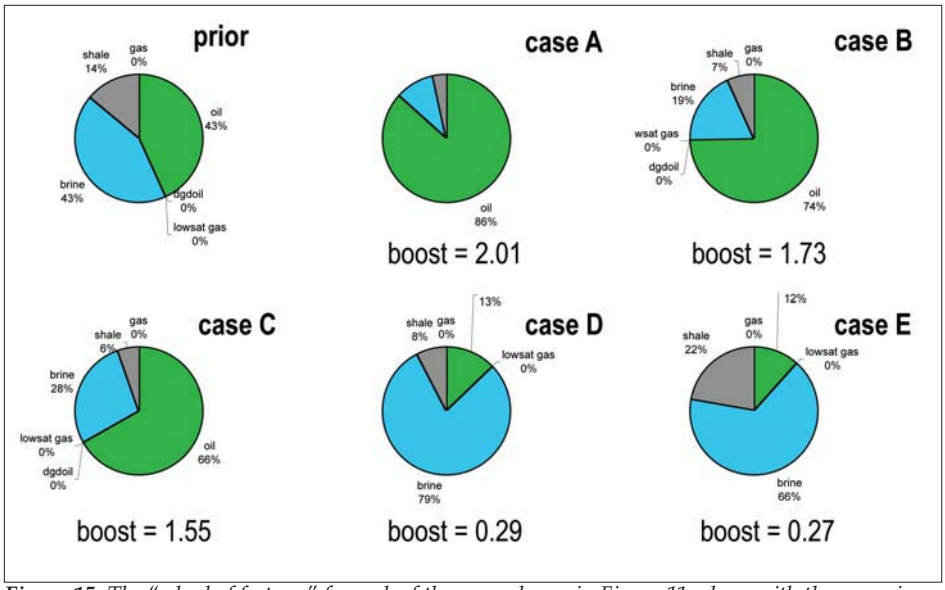

Figure 15. *The "wheel of fortune" for each of the cases shown in Figure 11, along with the one prior to taking the observed seismic into account.*

is not large enough to see a hydrocarbon effect. We will show that the situation is much more pessimistic than this conventional wisdom would suggest. A significant portion of the pessimism will come from the analysis of the differential amplitude. A 25-Hz Ricker wavelet was used to generate the synthetics. The noise level (standard deviation) is assumed to be about one third the size of the reflection in case D.

A Bayesian model-based inversion is done that adjusts the model parameters so that the synthetic of the model matches the seismic to within the estimated noise level. This is demonstrated by Figure 12 that shows the synthetic seismic of the ensemble of models before, then after inversion. The consequence of doing the inversion on the primary parameters (fluid type and N/G) of the model is shown in Figure 13. Note that the net sand, as well as the oil probability, decreases going from case A to case E. The quantitative effect on the net sand distribution in shown in Figure 14.

To calculate the effect on the "wheel of fortune" (the fluid probabilities), we have assumed a seismic confidence factor of 68%—a typical number for well-processed 3D seismic data. This means that 32% of the time there is a significant flaw in the data (or there is a scenario that has not been taken into account), and the seismic amplitudes should be ignored. For cases A–C, we assume that there is a structural change in the amplitudes consistent with a gravity effect on the less dense hydrocarbon. A further Bayesian boost consistent with increasing a 50% probability of hydrocarbon to 60% is applied. For cases D and E, the opposite is assumed, and a Bayesian boost of 40% is applied. The resulting changes to the "wheels of fortune" can be seen in Figure 15. The change to the oil probability ranges from an increase of 100% to a decrease of 70%. Detailed sensitivity analysis shows that the fluid probabilities along with the net sand are most sensitive to the seismic response, followed by the N/G, then the acoustic impedance of the bounding shale. Surprisingly little sensitivity is shown to the porosity of the sand.

Business value of using relative amplitudes. We will examine two ways that the differential seismic amplitude measurements can lead to business value. The first is the simple situation of assuming that one has the option to drill a prospective hydrocarbon resource and an option to do the relative amplitude analysis (one already has the conventional absolute amplitude analysis). The question that we need to answer is the value of doing the more reliable relative amplitude analysis. The second situation is a more complicated one of two parties bidding against one another for the option to drill the prospective hydrocarbon resource. One party does not have the better analysis and the other party does. The party with the better analysis knows that the other party does not have the analysis, while the party without the better analysis does not know that the other party does. While there is significant value in both cases, we will find that the value in the second is more than twice as much.

Let us outline our assumptions. The cost of doing the better analysis is US$100 000. The cost of drilling the well to test the opportunity is $100 million. Without any of the seismic amplitude information, the chance of a hydrocarbon resource being economic is 50%. Given that the hydrocarbon resource would be economic, the probability that the more reliable relative amplitude information would have a positive indication is 85%. Without the relative amplitude information (that is, with only the absolute information), this probability would be only 67%. Negative information would be equally pessimistic, with probabilities of 16% and 29% respectively. We also assume a 50% correlation between the less reliable and more reliable information (this will be needed for the competitive bidding situation). These reliabilities were determined by the size of the Bayesian update to the oil probability given in the last section (cases B and D).

Another important difference between the two types of information is the difference in their estimation of the net sand. The more reliable relative information will estimate more net sand given the more positive amplitude information than the less reliable information. The converse will be true when there is negative information—the more reliable information will indicate less sand. Again using the results of the last section, the value of the opportunity given positive reliable information and a positive drilling result would be $700 million, but with negative reliable information it would be only $400 million. Given less reliable information the value would be $600 million and $500 million (respectively, for the case of positive and negative less reliable information). We assume, in the competitive bidding situation, that the bidder with the less reliable information will bid 50% of the naive expanded net present value (ENPV) (that is assuming there is no bias to his winning the bid caused by the party with the more reliable information). The bidder with the more reliable information will bid 100% of the naive ENPV of the bidder with less reliable information.

Using a standard decision-tree-based, value-of-information methodology (Bickel et al., 2006), it can be shown that the value of the better information is $60 million for the first situation. This decision tree consists of two decisions and two uncertainties. The first decision is choosing the type of data, followed by the uncertainty of what those data will indicate.

Based on those data the decision will then need to be made whether to drill the prospect, followed by the uncertainty of whether the well is an economic discovery.

The value of the information in the second situation, competitive bidding, is $130 million. It makes the difference between losing $50 million if one bids without the more reliable information to making $80 million if one bids with the more reliable information. This is a classic case of "winner's curse," where the person with the less reliable information wins the bid only if the more reliable information indicates that it is not likely to be successful, and that if it is successful, it will not have as much value. This analysis consists of two decision trees—one for each of the two bidders. Both trees have three uncertainties and one decision. The first uncertainty is what the data will indicate. This is followed by the decision of whether to bid or not. The last two uncertainties are whether or not the bid is the winning bid and, if it is, whether the well is an economic discovery.

A further analysis of the competitive bidding situation highlights an interesting effect of having two parties bidding based on information with different reliability—the entity receiving the proceeds of the bidding will have a larger expected return. The expected return for the case above is $192 million. If both parties would have reliable information, the expected return would be reduced 15% to $162 million.

Conclusions. The use of relative seismic amplitudes has significant advantages. It mitigates the risk of not finding hydrocarbons and has a significant influence on the uncertainty of the net sand (thereby the NPV of the project, if hydrocarbons are found). For the case that we examined there was little sensitivity to the porosity and only mild sensitivity to the N/G. This can be specific to the situation, and the sensitivities should always be modeled for the case of interest. Finally, there is significant value to the application of this methodology, in the tens of millions of dollars per application. The value is amplified in competitive bidding situations.

Suggested reading. "Lithology and fluid prediction in lightly explored basins" by Duncan et al. (ASEG 2004 *Expanded Abstracts*). "Integration of uncertain subsurface information into multiple reservoir simulation models" by Glinsky et al. (*TLE*, 2005). "Application of integrated risking on a South African prospect" by Glinsky et al. (EAGE 2004 *Extended Abstracts*). "Delivery: an open source model-based Bayesian seismic inversion program" by Gunning and Glinsky (*Computers and Geosciences*, 2004). "Lithology and fluid prediction in lightly explored basins" by Duncan et al. (ASEG 2004 *Expanded Abstracts*). "Integration of uncertain subsurface information into multiple reservoir simulation models" by Glinsky et al. (*TLE*, 2005). *Geostatistical Reservoir Modeling* by Deutsch (Oxford University Press, 2002). "Stochastic reservoir characterization using prestack seismic data" by Eidsvik et al. (GEOPHYSICS, 2004). "Rapid spatially coupled AVO" by Buland et al. (GEOPHYSICS, 2003). "Quantifying 3D land seismic reliability and value" by Bickel et al. (SPE Annual Technical Conference, 2006). TLE

Acknowledgments: The authors would like to acknowledge the financial support of the BHP Billiton technology program, and to thank BHP Billiton for permission to publish the results. They would also like to thank Chris Haase for useful discussions.

Corresponding author: mglinsky@bhpbilliton.com