Deep Water Oil Sequential Exploration Test Study

The March 2006 issue of the journal Decision Analysis[1] (Volume 3, Number 1) contains a paper written by Eric Bickel and James Smith titled, Optimal Sequential Exploration: A Binary Learning Model[2]. The paper discusses a deep-water oil and gas exploration problem with incomplete information.

Using the oil exploration problem, we'll show how easy, fast, and natural it is to build models that contain incomplete information using the Provisdom Decision Platform. In addition to the joint probability distribution shown in the Bickel and Smith paper, we'll show how other incomplete information sets can be handled, including joint discrete and continuous-distribution information. We'll discuss how the solution to general incomplete information sets is determined by how the model builder chose to classify uncertainties. As always, we'll also find the strategy that maximizes shareholder value.

We'll begin by briefly describing the oil exploration project and the method used by Bickel and Smith to solve the problem. Then we'll discuss how the problem is modeled using the Provisdom Decision Platform to match Bickel and Smith's result. Next, we'll make the problem a bit more realistic by having the payoffs occur at the right time and lowering the risk-free rate. Finally, we'll imagine that more information is available and continue to add realism piece by piece, each building on top of the last:

- Include varying costs of moving the drilling rig from well to well
- Add an additional complex conditional probability
- Remove one of the unconditioned probabilities
- Include the uncertainty surrounding the oil volume of each well with oil and the oil price.

Oil Exploration Project

Bickel and Smith describe a corporation in the oil and gas industry that has six oil and gas wells numbered 1 through 6 that are geologically dependent. The corporation can drill a well to determine whether it has oil or is dry, but they can only drill one well at a time and each drilling takes six weeks. For each well, the corporation has information regarding the unconditioned chance of the well having oil, the NPV for the corporation given that the well has oil, and the NPV given that it's dry:

Well	Chance of Oil	NPV of Oil	NPV of Dry
1	35%	$60M	-$35M
2	49%	$15M	-$20M
3	53%	$30M	-$35M
4	83%	$5M	-$40M
5	33%	$40M	-$20M
6	18%	$80M	-$20M

[1] http://www.informs.org/site/DA/
[2] http://faculty.fuqua.duke.edu/~jes9/bio/Optimal_Sequential_Exploration.pdf

For most of the wells, the corporation also has information regarding the chance of a well having oil conditioned on one of the other wells, J, having oil:

	J = 1	J = 2	J = 3	J = 4	J = 5
P(2 oil|J oil)	59%				
P(3 oil|J oil)	63%	65%			
P(4 oil|J oil)	83%	83%	83%		
P(5 oil|J oil)	39%	55%	42%	33%	
P(6 oil|J oil)	31%	24%	31%	18%	26%

Since there is incomplete information regarding the conditional well scenario probabilities, Bickel and Smith minimized the mutual information between the well scenarios to arrive at a full joint probability distribution. Next they used dynamic programming with a six-week periodic discount rate of 1% (equivalent to an 8.653% continuously-compounded yearly risk-free rate) to solve for an optimal strategy and an NPV of $14.4M. Bickel and Smith point out that each well has a negative intrinsic value so the strategy of looking at each well in isolation would have an NPV of $0. Moreover, they point out that their strategy eludes the following rule-of-thumb strategies:

- Drill the well with the highest probability of success
- Drill the cheapest well
- Drill the well with the smallest expected loss
- Drill the well with the largest intrinsic value
- Drill the potentially largest well

The optimal strategy is shown below:

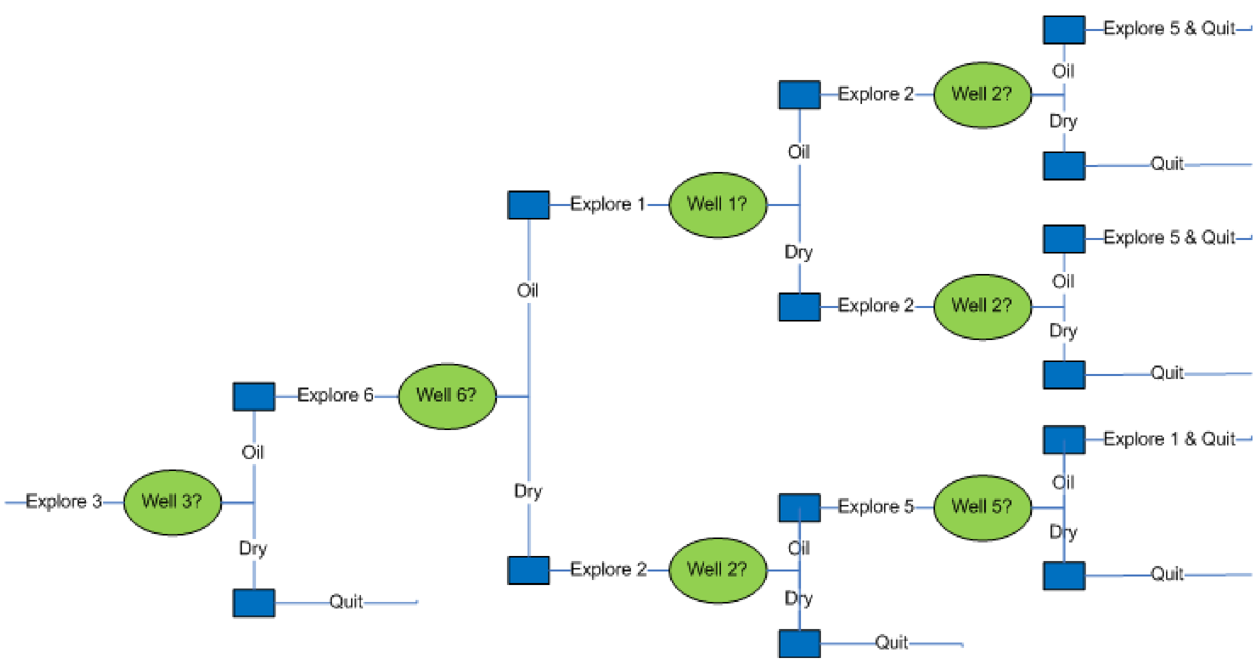

The relative frequency distribution for exploring each well is shown below. This graph provides a nice overview of the strategy. Exploring Well 3 has a relative frequency of 100%, as it is explored on all paths, while the relative frequency of exploring Well 4 is zero.

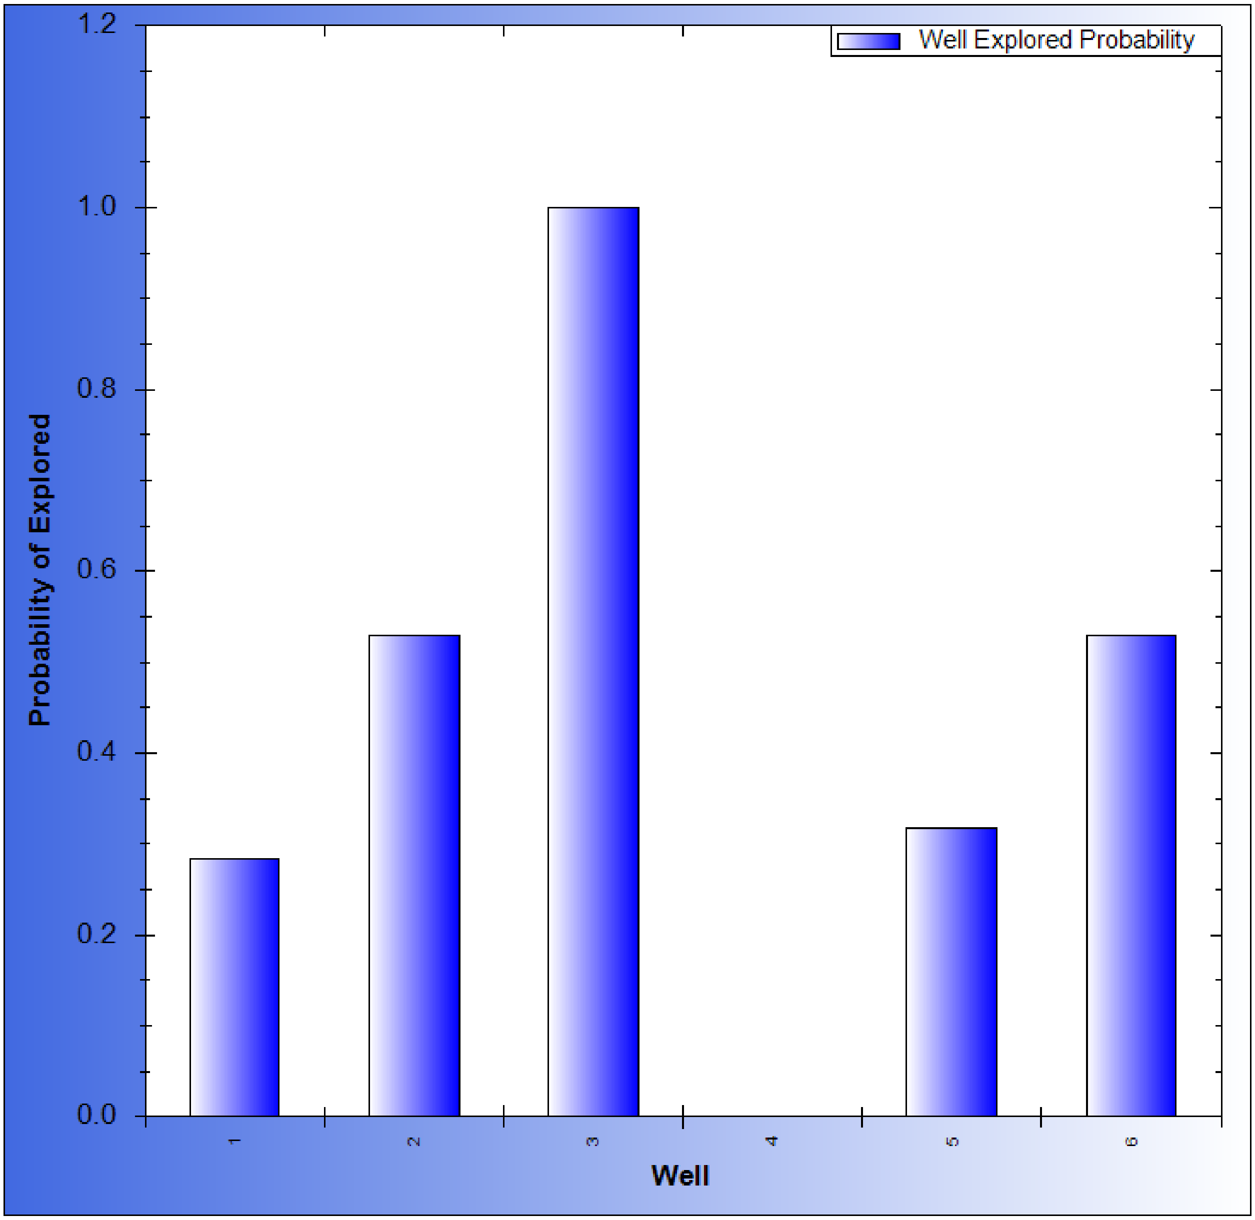

Building the Sequential Exploration Problem

The oil exploration problem can be represented by two basic information rules:

1. Choose whether to drill and which well to drill each period
2. Resolve whether the chosen well has oil or is dry

Each well has an associated scenario for whether it has “Oil” or is “Dry”. The unconditional probabilities of “Oil” as well as the pairwise conditional probabilities are entered as constraints on the distribution. The full distribution is found numerically using Provisdom’s information theoretic framework.

There are three state variables:

1. Time
2. List of current states of the wells (unknown, oil, or dry)
3. Current well being drilled

Both the strategy and the NPV from the Bickel and Smith paper were matched.

Payoffs Occur after Well Scenarios are Resolved

In the original model, the payoffs for the wells having oil or being dry occurred one period before the resolution of the uncertainty. We built a model in which the conditional payoffs occurred after it was known whether the well had oil or was dry. This lowered the NPV from $14.403M to $14.260M.

Lowering the Risk-free Rate

Since a risk-free rate of 8.653% seems unrealistic in 2007, we lowered the rate to 5%. This change increases the NPV from $14.260M to $14.454M.

A portion of the solution is shown below.

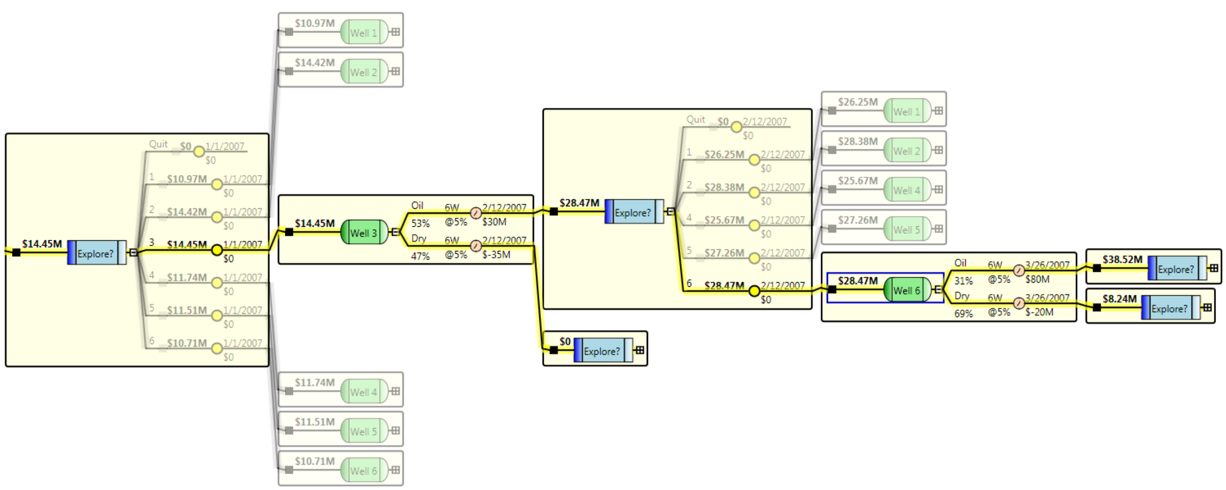

Finally, let’s examine the gradient of the NPV with respect to the unconditional probabilities:

Well	Gradient ($/%)
1	385947
2	248050
3	621619
4	-184
5	-294922
6	-255622

The gradient values tell us the numerical derivative of the NPV with respect to each quantity. We see that wells 1-3 show a positive relationship between their unconditioned probability of having oil and the NPV, while wells 4-6 have a negative relationship. To see why lowering the unconditional probability could actually increase shareholder value, consider Well 4. The unconditional probability of Well 4 having oil is equal to the probability of Well 4 having oil conditioned on wells 1-3 having oil. If we lower the unconditioned probability of Well 4 having oil without also lowering the conditional probabilities, then the probability of Well 4 having oil increases when any of wells 1-3 have oil. Therefore, we have increased the positive relationship between wells 1-3, which increases the NPV. For example, lowering the unconditional probability of Well 4 having oil from 83% to 75% increases the NPV from $14.454M to $14.758M. Of course, if we had kept the original belief that Well 4 is independent of the other wells, the gradient for Well 4 would have been zero since it is never drilled.

Model Extensions

Cost of Moving Drilling Rig

Now we'll imagine that it costs to move the drilling rig from well to well (thanks to Jim Smith for the idea). The cost of moving the drilling rig between wells is shown in the following table:

	Well 2	Well 3	Well 4	Well 5	Well 6
Well 1	$1M	$3M	$8M	$4M	$6M
Well 2		$2M	$9M	$3M	$5M
Well 3			$6M	$1M	$5M
Well 4				$4M	$6M
Well 5					$1M

Including these costs changes the optimal strategy substantially and lowers the NPV from $14.454M to $11.918M:

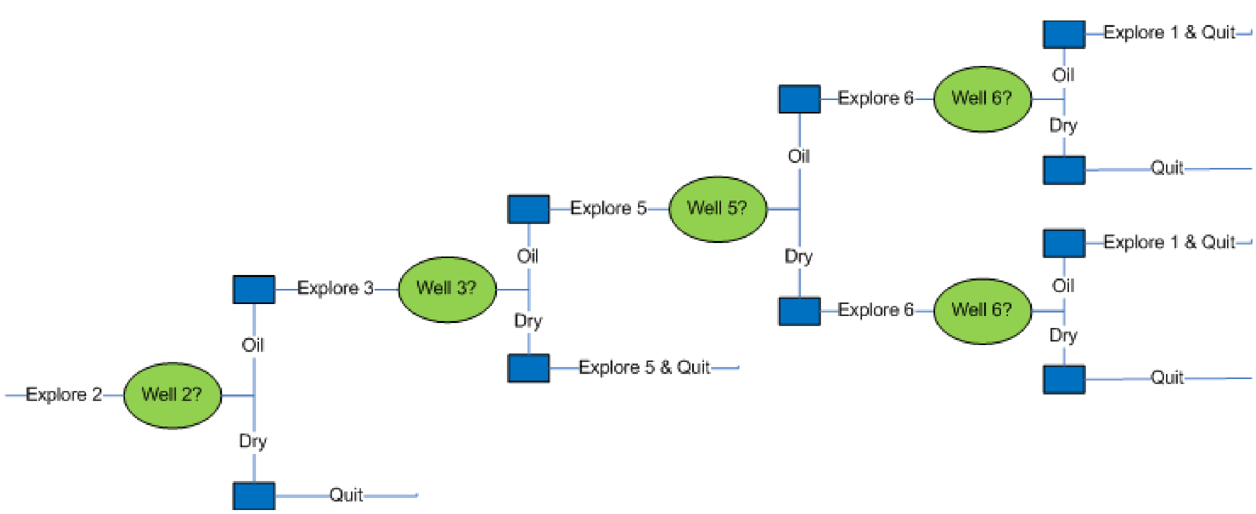

Adding a Complex Conditional Probability

We'd also like to explore how the optimal strategy would change with one more piece information. We added that the chance of Well 1 and Well 2 having oil given that Well 2 has oil and Well 4 has oil is 30%, i.e. P(1 oil & 2 oil|3 oil & 4 oil)=35%. This one added piece of information increases the NPV from $11.918M to $12.712M and again changes the optimal strategy:

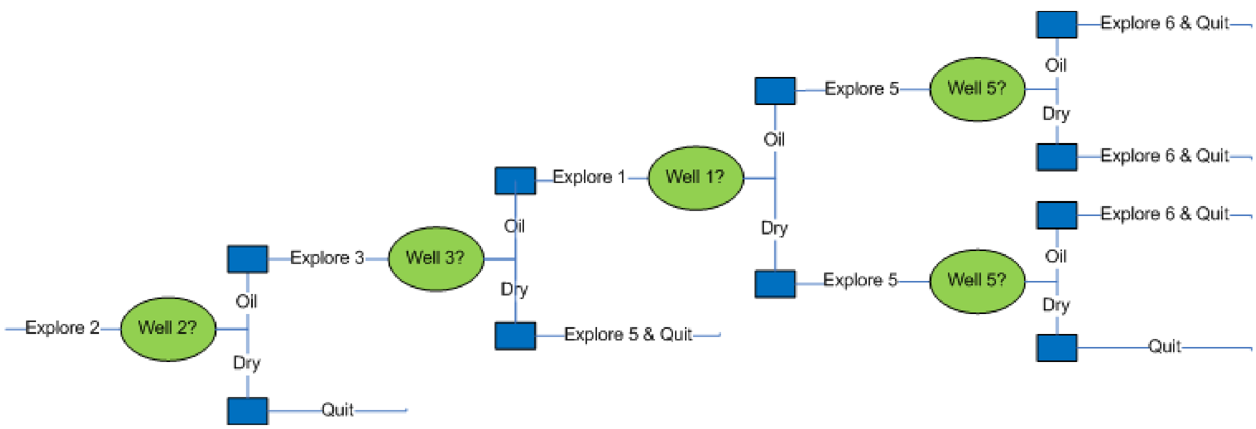

Removing the Unconditional Probability of Well 1

Now suppose we remove the unconditioned probability that Well 1 has oil. This changes the probability of Well 1 having oil to 42.3%, the NPV from $12.712M to $15.876M, and the optimal strategy to:

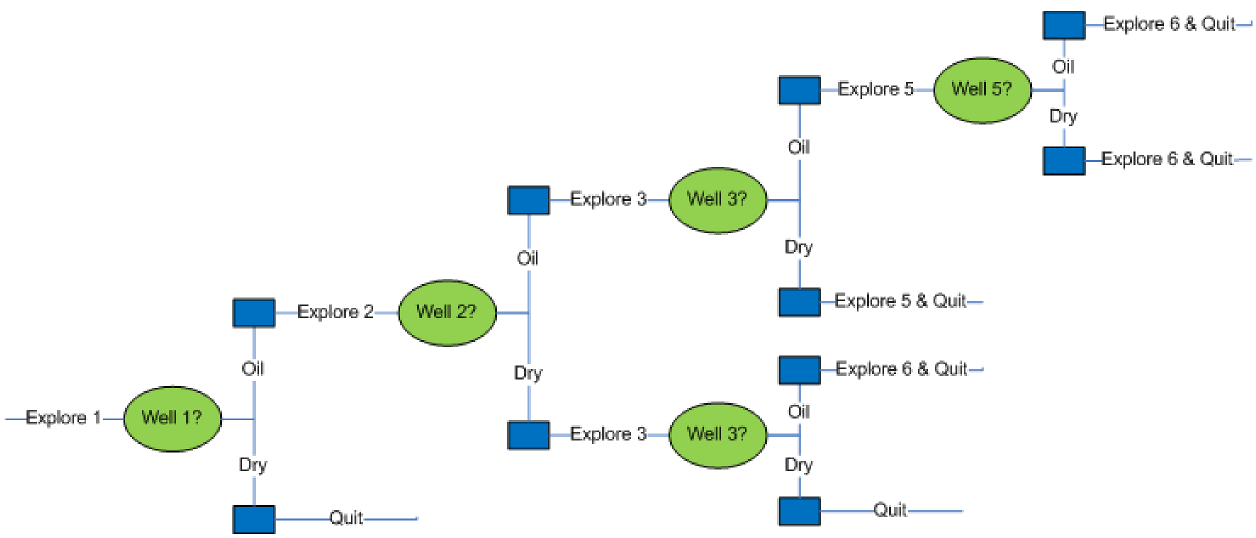

The relative frequency of exploring each well is shown below:

Well	Relative Frequency of Explored
1	100%
2	42.3%
3	42.3%
4	0%
5	25%
6	26.7%

Bickel and Smith note that for the original set of information, minimizing mutual information is equivalent to maximizing entropy. By default, Provisdom's system neither simply maximizes entropy nor minimizes mutual information, but instead realizes that a model-builder chooses uncertainties in a way to maximize the relevant information going into the model.

With the removal of the unconditioned probability, the equivalence between maximizing entropy and minimizing mutual information no longer holds. By setting the Qualitative model property to "Arbitrary", we can investigate how the solution would change if we had maximized entropy (this would imply that the decomposition of uncertainties in this model had been chosen arbitrarily, which of course they were not). This would have had the same optimal strategy as the "Maximum Information" solution but would have changed the probability of Well 1 having oil to 43.6% and increased the NPV to $17.415M. Even though the optimal strategy didn't change, the probabilities throughout the model did. The well-explored relative frequencies are now:

Well	Relative Frequency of Explored
1	100%
2	43.6%
3	43.6%
4	0%
5	25.7%
6	27.5%

Alternatively, we could have set the Qualitative parameter to "Minimize Interactions" (implying that there is reason to believe there should be a minimum of connections between the uncertainties), which would have changed the probability of Well 1 having oil to 40.4% and lowered the NPV to $14.643M, again showing that information interactions between the wells is providing value to the corporation's shareholders. In that case:

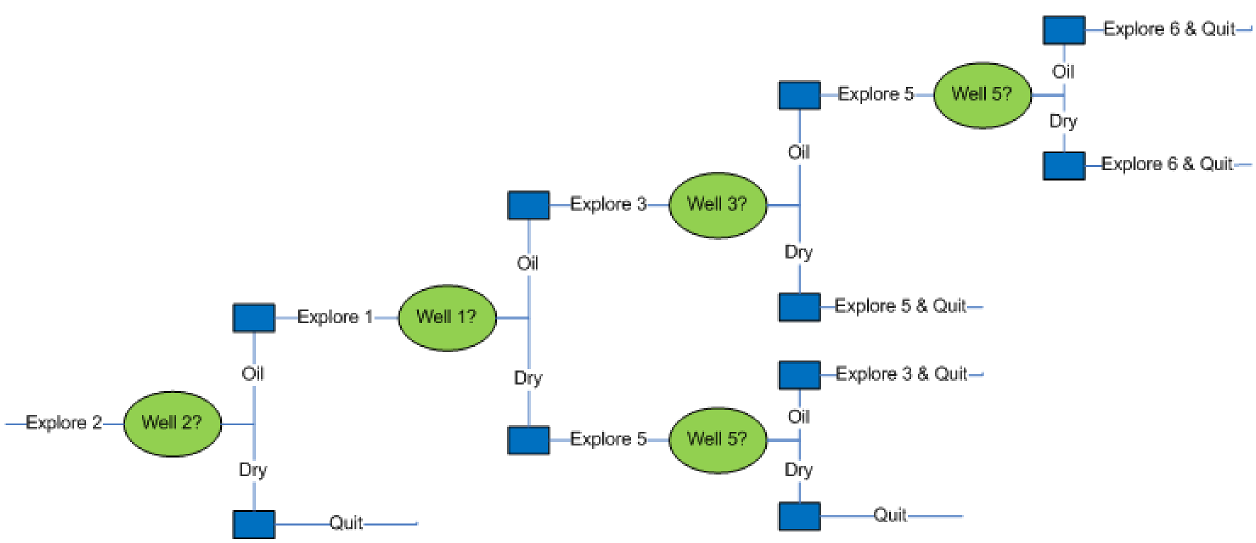

The well-explored relative frequencies are:

Well	Relative Frequency of Explored
1	49%
2	100%
3	37.1%
4	0%
5	49%
6	17.8%

Note that minimizing mutual information is not equivalent to "Minimize Interactions" and "Arbitrary" is only equivalent to maximizing entropy when the model contains no uncertainties other than scenarios.

Oil Volume and Price

Now we'll imagine that the oil volume and oil price is uncertain. Due to the large number of possible states, Provisdom's current engine on an off-the-shelf laptop takes just under two minutes to run this model with four wells. Suppose the current price per barrel of oil is $100, and that the corporation's information tells them that oil price has an expected growth rate of 3%, a yield of 2%, and a volatility of 40%. The corporation's information determines the oil volume standard deviation for each well. To closely match the original information, we set the drilling costs equal to the 'dry' payoff from the original information and the oil volume mean equal to the difference between the 'oil' payoff and the 'dry' payoff and divided by $100 (current price of oil):

Well	Drilling Cost	Oil Volume Mean	Oil Volume Standard Deviation
1	$35M	950K barrels	300K barrels
2	$20M	350K barrels	200K barrels
3	$35M	650K barrels	300K barrels
4	$40M	450K barrels	200K barrels

The corporation also has additional various types of information regarding the relationships between oil volume and chance of having oil:

- There is a 50% correlation (log) between Well 1 Oil Volume and Well 2 Oil Volume.
- If Well 2 has oil, the Well 3 Oil Volume mean increases from 650K barrels to 700K barrels.
- If Well 4 Oil Volume is 500K barrels, the probability of Well 3 having oil increases from 53% to 58%.

There is incomplete information regarding the conditional probabilities of having oil and conditional distributions of oil volume, but the Provisdom Decision Platform can solve for a joint distribution and maximize shareholder value given the information that we do have. The NPV is $39.247M and the optimal strategy is to begin by drilling Well 1. If Well 1 is dry and the price of oil falls, quit. Otherwise, continue by drilling Well 2. The rest of the strategy is too complex to describe in detail here.

A portion of the solution is shown below.

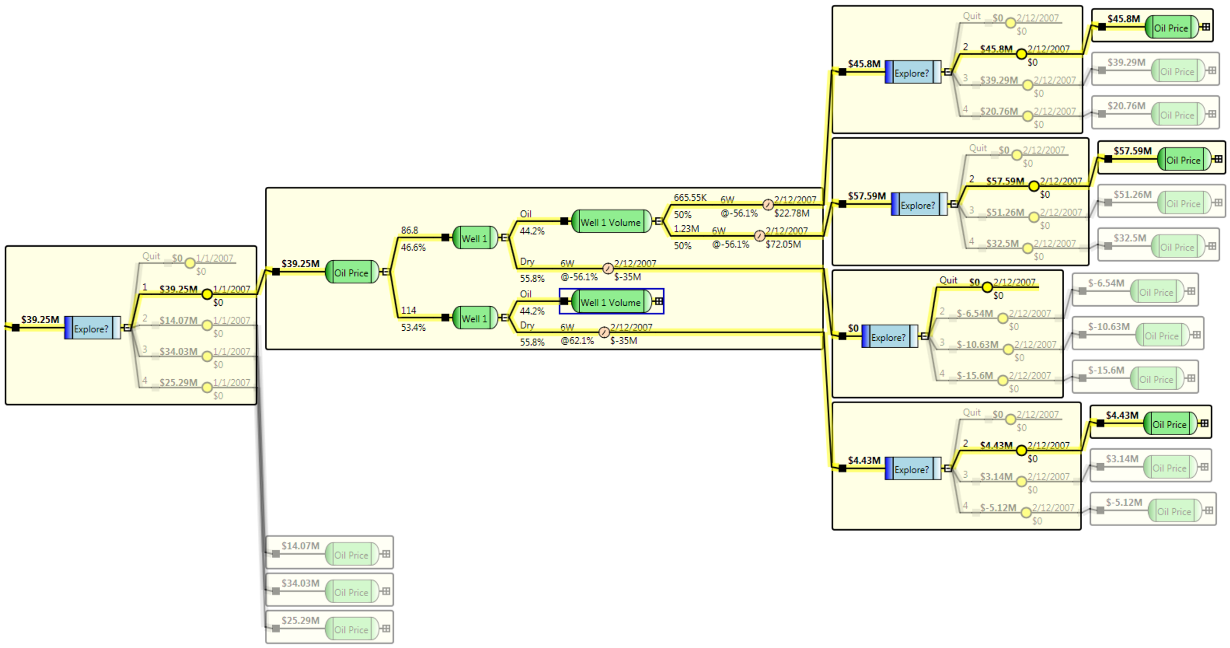

Below is shown the NPV as a function of starting oil price.

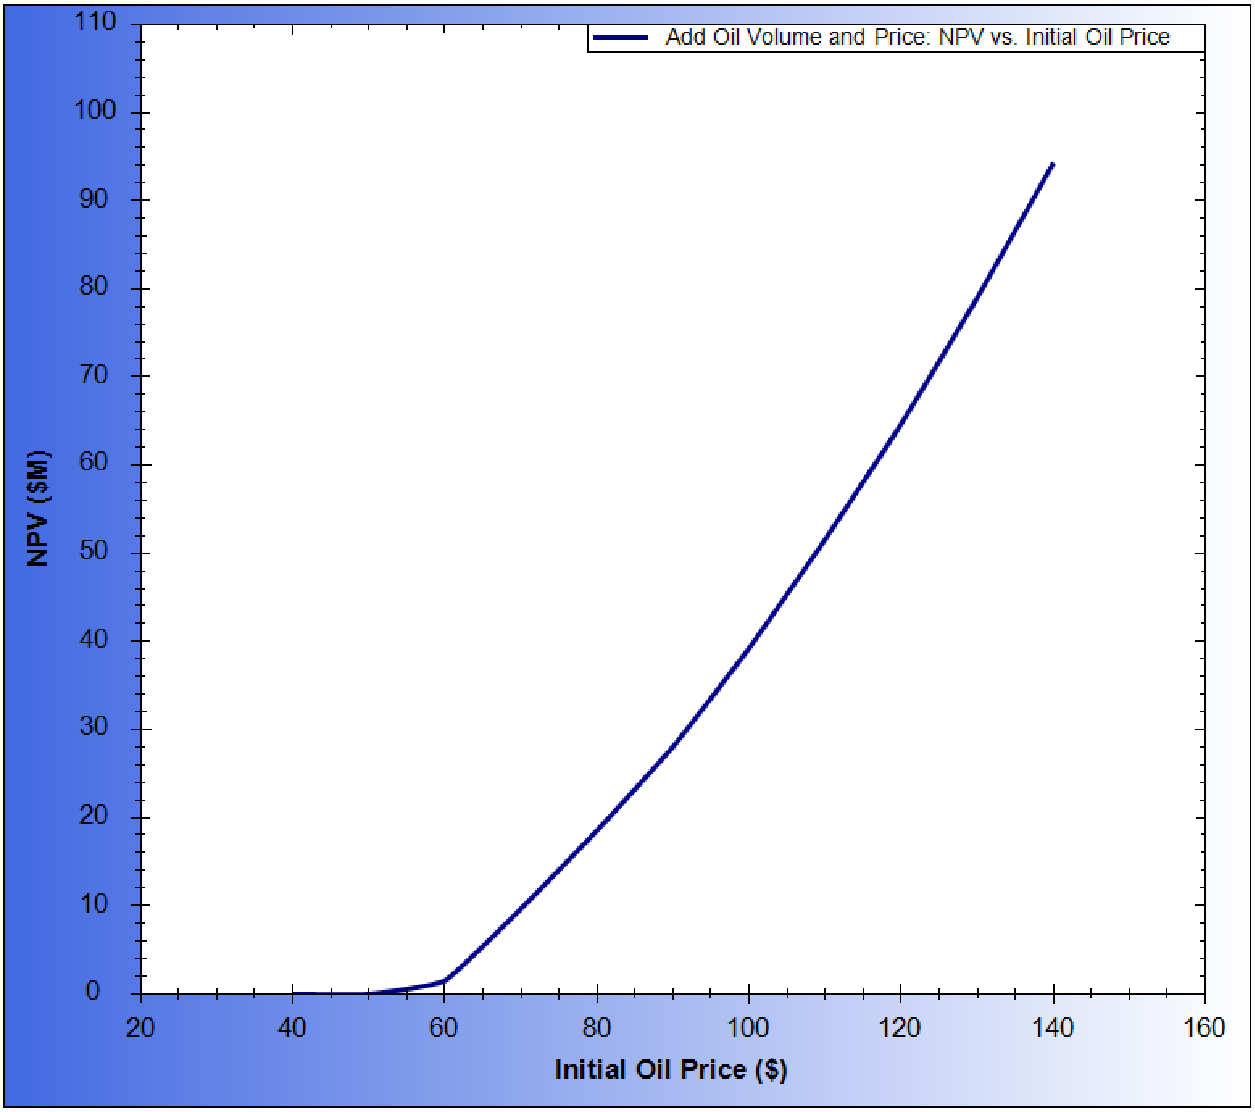

We can do a search query to find the initial oil price where the strategy changes. This occurs at about $58.10, at which point our initial choice switches from "Quit" to exploration of Well 1.

Discussion

Modeling the original oil exploration problem with the Provisdom Decision Platform simply required creating two basic Information Rules for each well, an associated scenario for each well, three state variables, and elucidating the model properties. Incomplete information was properly handled in a number of cases, thereby removing the burden of the model-builder to include information they don't have and increasing the accuracy of the model. The oil price was accurately modeled as an asset and the discount rates changed accordingly. The base model took about an hour to build and the addition of uncertainties for oil volume and oil price required about an additional hour to add to the model. By simply providing information that a petroleum corporation already has available, in a few hours the corporation could potentially create an improved drilling strategy using the platform that generates an additional $40M in shareholder value when compared to the default strategy of not drilling wells with a negative intrinsic value.