BP's Deepwater Oil Spill: A Statistical Analysis of How Many Relief Wells Are Needed

Posted by JoulesBurn on June 14, 2010 - 8:25am
Topic: Environment/Sustainability
Tags: bp, gulf of mexico, oil spill, original, relief wells [list all tags]

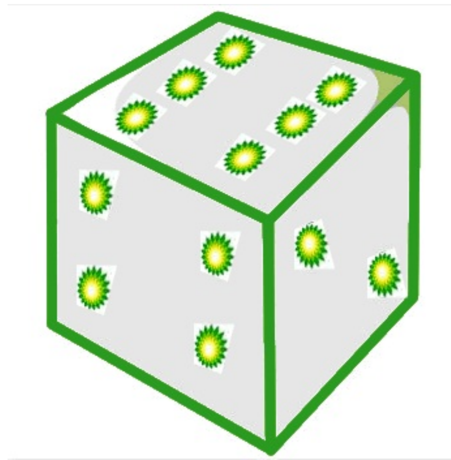

The planned permanent solution for the BP Macondo spill in the Gulf of Mexico is the killing of the well by drilling a relief well to intersect the blowout wellbore, enabling the placement of high density mud and eventually cement. Currently, two relief wells are being drilled, one started on May 2 and the second on May 14 (although the latter was delayed a few days later on). On The Oil Drum and elsewhere, many have been questioning why only two relief wells have been sunk, given the risky and uncertain nature of the kill process and the long time lag in getting more wells drilled if the first two are unsuccessful. There are many technical, political, and economic arguments that can be used to justify the need for more wells. What I will do herein is develop a statistical model which can be used to weigh the potential benefits of additional wells added this late in the crisis. One of the more critical factors is time -- the time it takes before the blowout wells is killed. Does drilling more relief wells change the expected time before the kill?

There is some recent news and concern about the relief wells and the delays in getting started:

> But BP didn't begin drilling the relief well until 12 days after the start of the disaster as the company and government rushed through environmental reviews, permits and other plans. The government does not require oil companies to have relief well plans in place ahead of time, and the lack of planning cost the company valuable time to get the spill under control.

The drilling seems to be on schedule,

> BP says the relief well has been a success and ahead of schedule, representing a welcome change for engineers who have been attempting one risky, untested maneuver after another. Relief wells are a more proven method in the industry, and engineers are comfortable and confident in the process.

but it still will be a while.

> Kent Wells, BP's senior vice president of exploration and production, said this week that more details would be released when the process nears completion in early August.

Should more have been drilled, based on the difficulties of the task and the chances for success? How do these chances affect the time frame we are looking at to get the well plugged? Here is one way of looking at the problem, starting with a game of chance.

Roll the Dice

Consider the following game which consists of the rolling of a single die. A single play consists of rolling it again and again until you get a six, counting the number of rolls it takes. Play the game many times, and keep score like this:

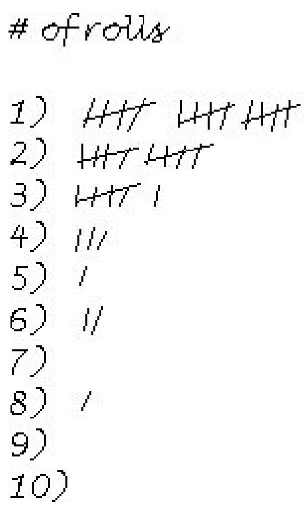

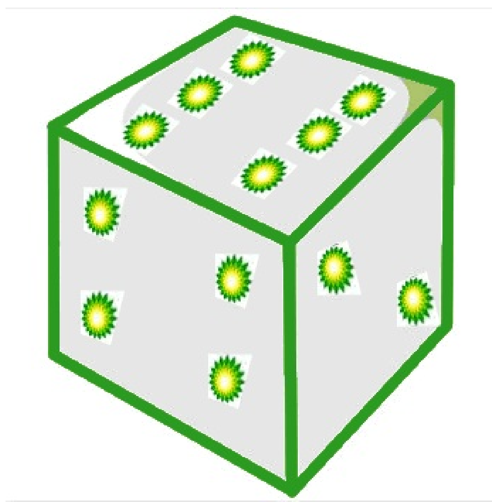

After many, many plays, the statistics get rather good and you can make a bar chart showing the frequency of each roll count divided by the total number of plays:

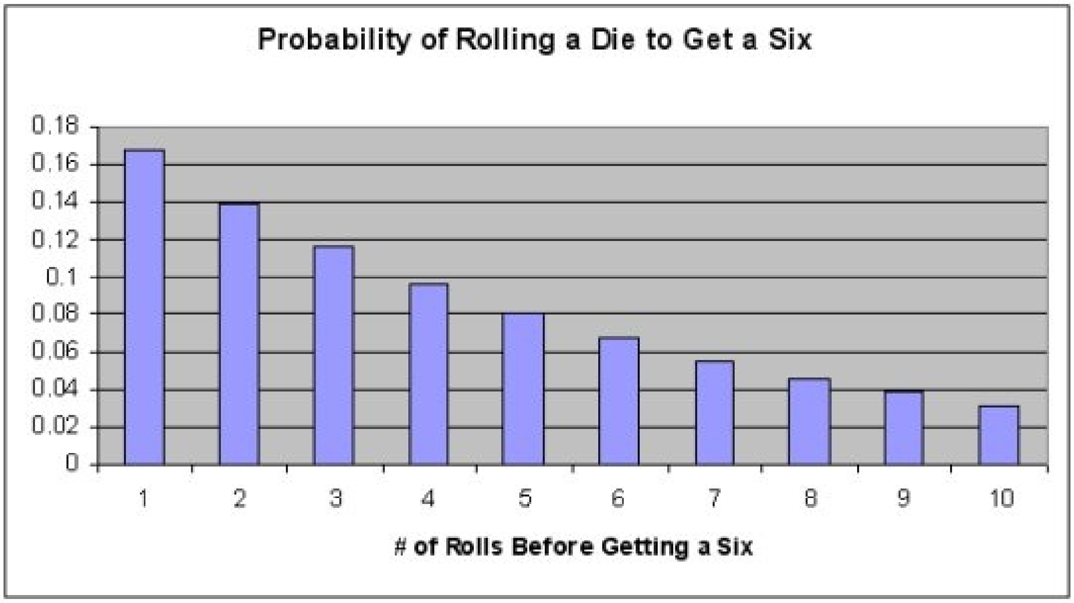

To be more exact, you would need to go higher than ten (infinitely high, actually), but you get the idea. This shows the statistics of how many rolls it takes to have a six come up. One is the most likely number of rolls, although it is of course more likely to take more than one.

Now, let's change the game a little. We secretly replace the die with one that explodes occasionally, ending that particular play. So you grab another (similarly loaded) die and play again -- but making sure to count the play that went awry as an attempt. The same probability profile would appear, except that the overall probability for getting a six would be less.

Now let's change the game a bit more. Instead of the number of rolls it takes, we are interested in the amount of time it takes to get a six. There is an average amount of time it takes to pick up the die and roll it again. If it took five rolls on one particular play, it would take something like five times the average for one roll. There would be some variability in the time, of course, and the spread would increase with the particular roll count we are interested in since there is a contribution from the spread of previous rolls. Before I get more specific on how this is done, or what happens if we roll more than one die, let us leave the analogy behind. But not before I add one more twist: the first roll seems to take forever.

Rolling Relief: The Assumptions

Let us consider the time probability for killing the blowout well with a relief well in the same way. There is an initial drilling period (rather long) to reach near the blowout wellbore, and then a certain probability that it will be successfully intersected, called P_s. If the try fails, the drillbit is backed up some, the hole is plugged, and a new attempt to find the wellbore is made. While not necessary, it will be assumed in this model that the probability for success remains the same for each retry. Also included is a probability that, during each additional attempt to intersect the wellbore, that the relief well becomes useless, perhaps due to a hopelessly stuck drillstring. I will call this the probability of "utter failure", or P_u.

The time needed to drill to the vicinity of the blowout wellbore will vary, as a number of things will affect the rate at which the drill bit moves. If one were to attempt the same task many times independently with the same starting conditions (obviously not possible), the set of the times required would take the form of some kind of statistical distribution. I will assume the normal

(Gaussian) distribution, as shown below (e.g. mean time of 70 days and a standard deviation of 5 days).

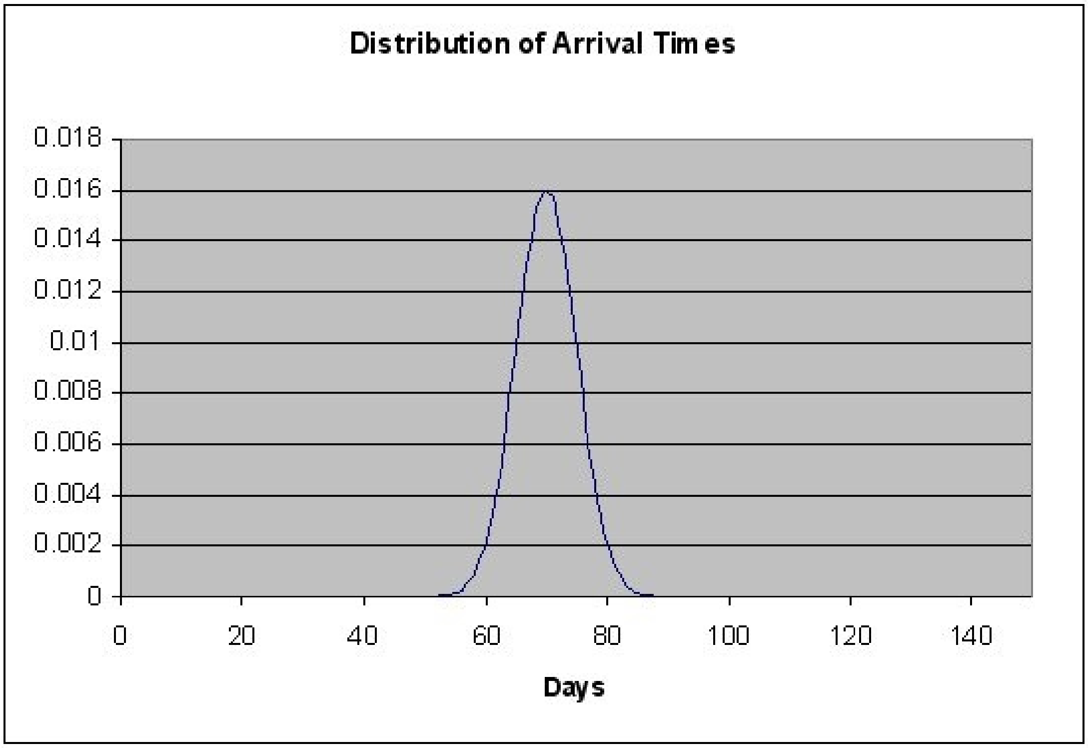

Let us assume a 20% probability for success (P_S=.20), which could include both hitting the wellbore and successfully filling it with mud and cement. Thus, the Gaussian above would be multiplied by that factor. This gives us the probability at any point in time. What we are more interested in, though, is how this probability is manifested in time; that is, how long do we wait for success. Instead of the probability distribution, we can look at the cumulative probability distribution, again scaled by the success probability:

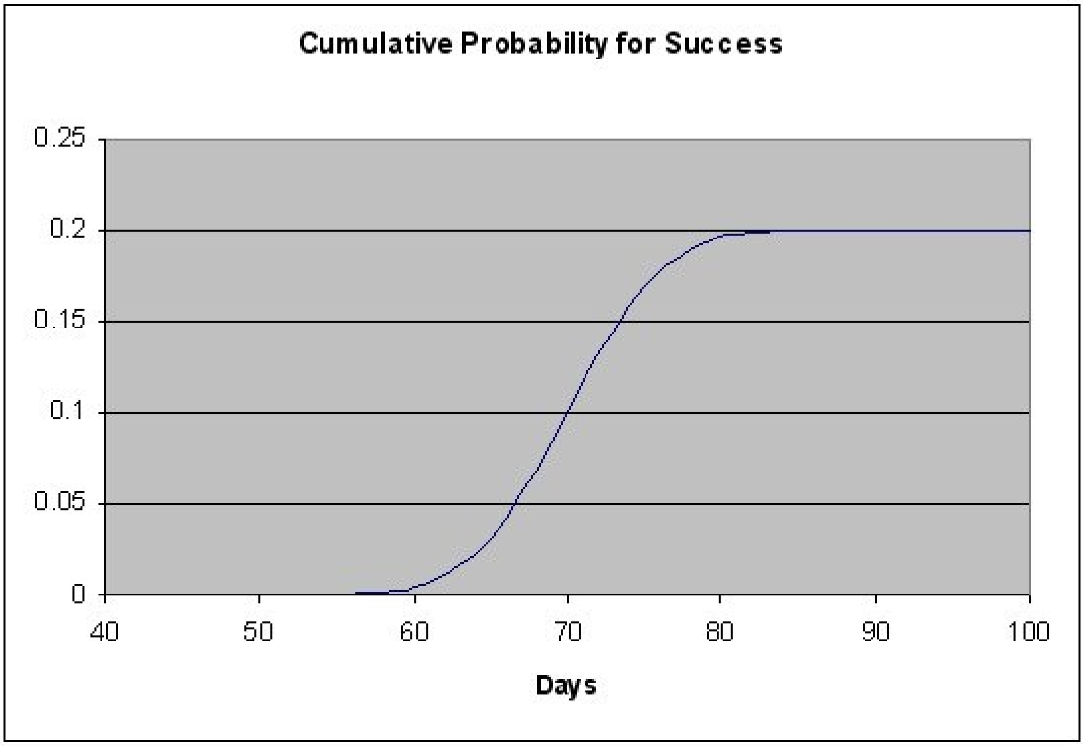

This is arrived at by calculating a running sum of the probabilities from left to right. With the values assumed in this example, it takes about 80 days for the full effect. However, since we only have a 20% success rate, this is not likely to be a successful resolution of the problem, and many wells would be required to insure a kill. For a given relief well, however, we get more than one crack at it. The drill bit is backed up, the hole is plugged, and a sidetrack is drilled for another attempt. This retry consumes additional time, and we can also describe this time delay with a Gaussian distribution with a separate mean and standard deviation.

The effect of the additional delay is as follows: if we consider the distribution as a large set of kill attempts in identical circumstances, 20% of these would be successful within about 80 days. Aside from a small fraction that will have failed completely, the remainder of the set will continue on with a retry. There will be a time delay, but we also have an additional spreading of the time distribution. This can be mathematically described as the convolution of the initial time distribution with a second distribution corresponding to the retry. Visually, one can describe various small segments (or discreet points) of the original distribution each giving rise to its own new distribution. The sum of all such distributions comprises the overall distribution of arrival times for the second attempt. This will be another Gaussian; conveniently, the convolution of two Gaussians is another Gaussian with the standard deviation equal to the sum of the variances of the two and with a mean equal to the sum of the two means. Defining the following parameters:

Mean Time for Initial Drill	t_1
Spread for Initial Drill Time	SD_1
Mean Time for Redrill	t_2
Spread for Redrill Time	SD_2

the successes from the initial attempt will be characterized by:

$G_1 = P_s \cdot \text{Gaussian}(t_1, SD_1)$

where G() is a normalized Gaussian distribution with the given mean time and SD. Those failed attempts that remain for each retry will give rise to a new Gaussian, displaced in time by that needed for the additional drilling etc., and this will have the form has the form:

G_i = Gaussian(t_2,SD_2) * Gaussian(t_{i-1},SD_{i-1})

Because of the convolution properties of Gaussians, this can be arrived at by constructing a new Gaussian displaced from the previous by the mean time for a retry and with a standard deviation somewhat larger:

SD_i = SQRT(SD_{i-1}^2 + SD_2^2)

and a scaling factor P_i which can be calculated for each retry recursively as:

$$P_i = P_s \cdot \sum_{j=1}^{i-1} \left(1 - P_j\right) - P_u$$

Alternately, one could for each retry perform a discreet convolution of the retry broadening/delay function with the previous result (if something besides Gaussians were used, for example). The back-up-and-redrill process can in principle be repeated infinitely many times, although there is probably a practical limit to how many side holes can be drilled in the single relief well.

Here is an example showing success probability distributions for a single relief well with multiple retries, with each retry contributing its own Gaussian:

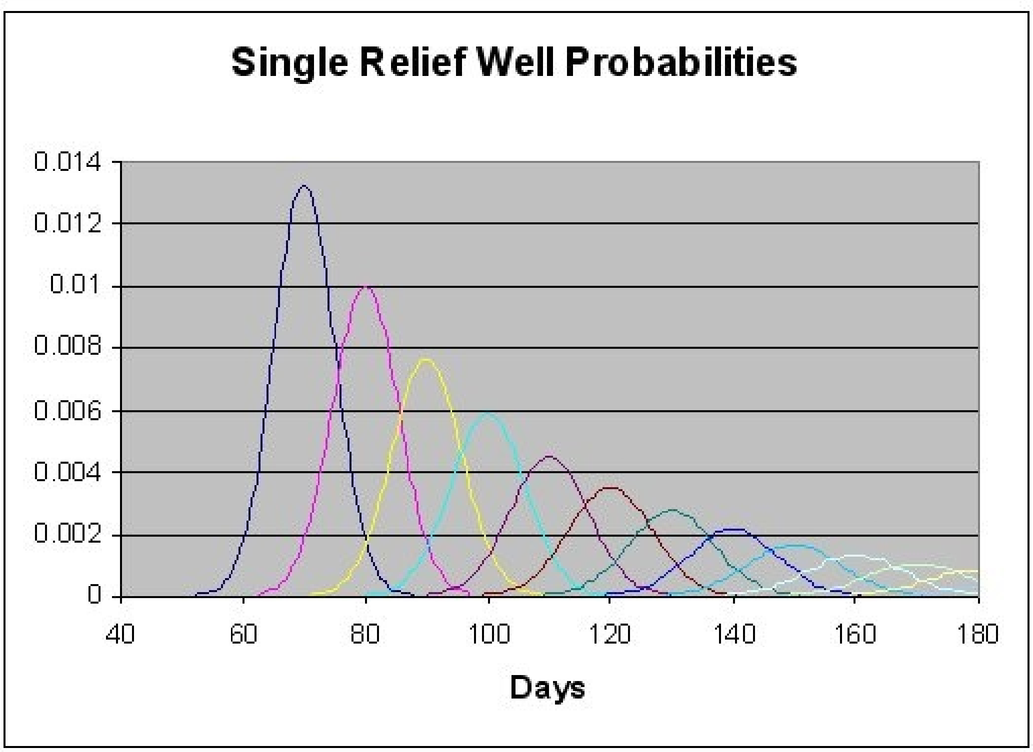

The following values are assumed:

Mean time to drill well (days)	70
Mean time for each re-attempt (days)	10
Standard deviaton for initial drill time	5
Standard deviaton for each re-attempt	2
Probability for intersecting well per try	0.2
Probability for loss of well per try	0.05

The overall probability distribution for success will be the sum of the individual Gaussians. To see how this impacts the anticipated time to kill the well, we look again instead at the cumulative distribution for all successes:

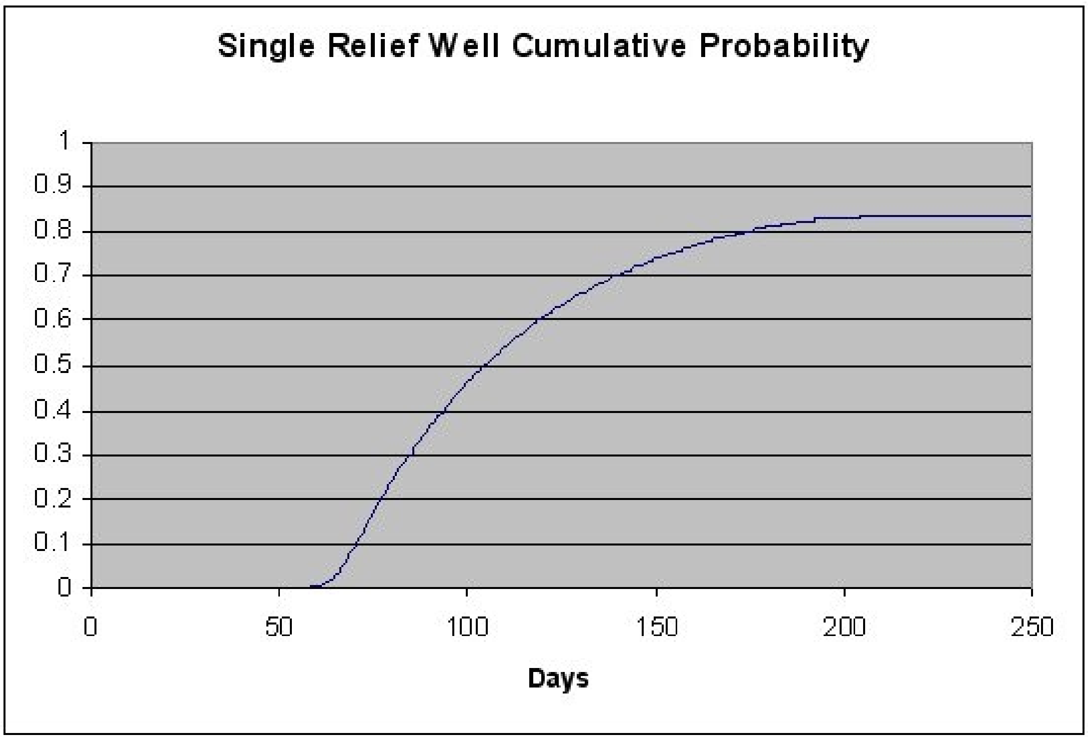

This illustrates the effect of the delay time for each retry, and why a single number, the overall success probability for a relief well (even if one can arrive at such a number) does not fully describe the situation. That the probability converges on a number less than unity is due to the small probability of "utter failure" with each attempt, meaning that the relief well is abandoned, and also because a finite number of retries were included. As it is, the overall success rate with the chosen parameters is higher than that occasionally mentioned for a single relief well (~75%). To get better odds -- and to decrease the amount of time before the blowout is killed, and hence the amount of oil potentially spilled -- more relief wells are needed.

More Wells = More Relief?

To summarize where we have come thus far, I started with one relief well with a modest (20%) probability of killing the well in one try. Because of the time needed to drill it, and some variability in the amount of time, this probability shows up within a small window of time (~70 days from the start of drilling). Next, I added the possibility of more than one try with this well (to back up a bit for another attempt takes some time). Because we take more swings, the eventual probability gets much higher (over 80%), but there is a time lag in getting there. So now we add more relief wells, identical to the first except perhaps with regards to the start time.

If we assume that additional relief wells have the same probability for success that the first had, it follows that each well will have the same cumulative probability distributions as the first. But how do these act in concert? A simplistic assumption is that they just get summed together. However, it is easy to see how this is not correct. For example, flipping two coins does not give 100% probability of one of them landing with heads showing. The solution is to consider what is required for the oil to be still flowing at a specific time: none of the relief wells will have worked by then. The probability of that being the case at a specific time, for each well, is 1 minus the cumulative probability distributions (CDP) at that time. Furthermore, the combined probability that no wells have succeeded is the product of these calculated values as follows:

$$CPD_{total}(t) = 1 - \prod_{wells} (1 - CPD(t))$$

Let us consider the cases of 1-4 relief wells started at the same time (increasing number of wells towards the upper left):

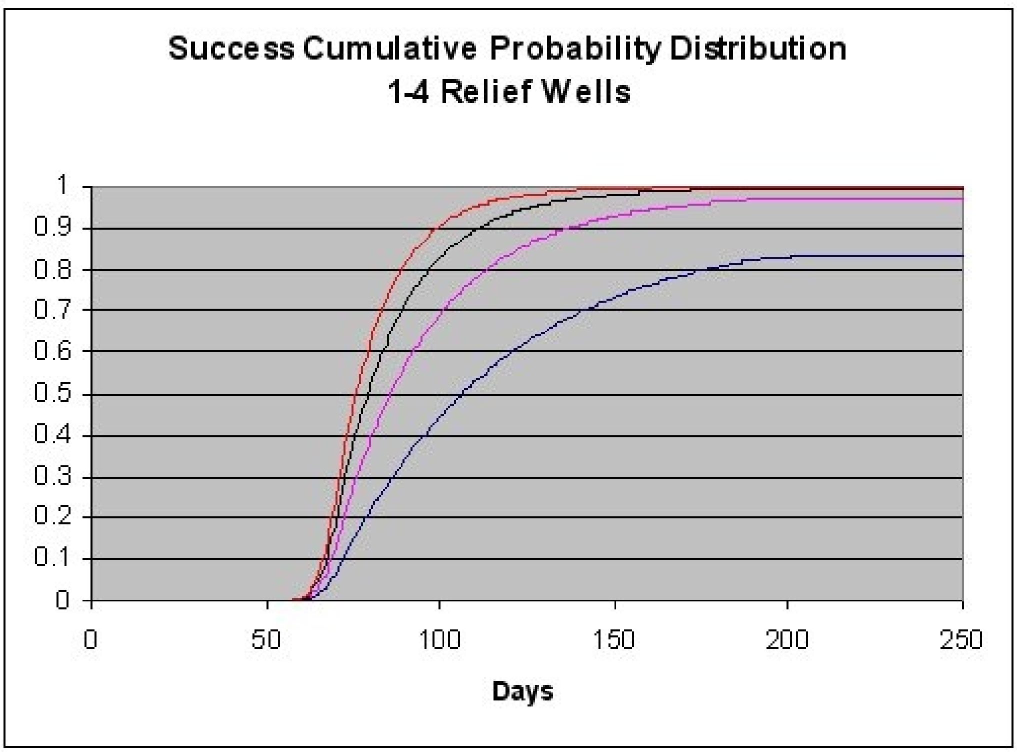

One effect is to increase the ultimate success probability, primarily after the addition of a second well. Another effect, though, is to decrease the time elapsed before a certain probability of success is reached. For example, it takes about 178 days to reach 80% probability with one well, 115 days with two wells, 96 days with three, and 87 days with four. With thousands of barrels of oil flowing into the sea per day, this is a lot of pollution potentially mitigated by a couple more relief wells.

Unfortunately, we no longer have the luxury of starting relief wells at time=0. The second relief well is perhaps 18 days behind the first, and no others are scheduled. Does adding more now make a difference? Here is the effect of adding wells 3 and 4, at +50 days, to the two currently being drilled at (wells 1 and 2 at 0 days and +18 days).

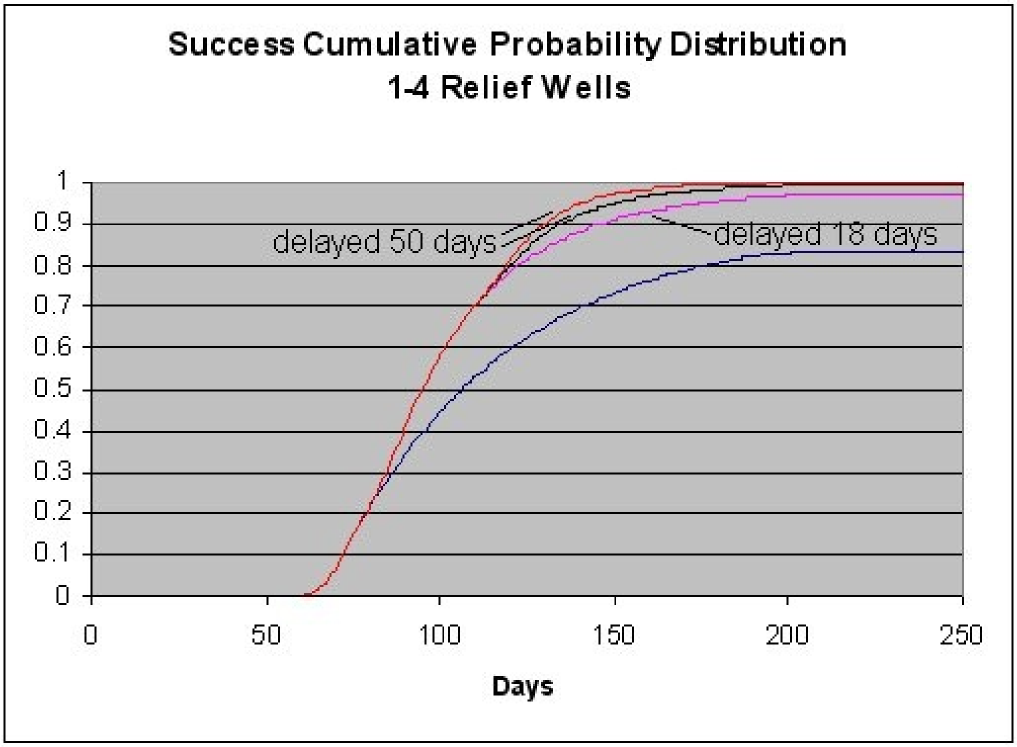

Although the increase in the ultimate success probability is the same, the benefit at the 80% probability level is minimal. Is it worth it, then, to drill more wells now? Certainly, it is logical to conclude that the benefit of additional wells at this point is much less than if they were started right away.

Should More Be Drilled Now?

Even if this model is a good one, the probabilities for relief well success are not known with any degree of confidence. I have chosen parameters which are consistent with predictions of when the well could be killed (August), the approximate probability of success, and time needed to regroup and try again after a swing and a miss. But the results suggest that, barring the loss of one or more of the current wells, additional wells will not significantly affect the time before the blowout is quenched. Readers are invited to download the spreadsheet used for this analysis, change the assumptions to those which you find more defensible, and offer up the results for further discussion.

In spite of that, Many will argue for at least one more well anyway, and I would probably fall into that camp. Wells are drilled all the time based on the chance for a big payout, and then never go into production. Hopefully another won't be needed, but the risk of not capping the well as soon as possible should be obvious to BP, and waiting until circumstances force the issue of another well will not go over well. In retrospect, they should have looked at the calendar (with the upcoming stormy season approaching), considered yet another unthinkable scenario regarding relief well success, and planned accordingly. In lieu of that happening several weeks ago, the time is now for BP to go above and beyond what they think is required. And with the deepwater drilling moratorium in place, many rigs are looking for something to do.

The spreadsheet which was used for these calculations can be downloaded here.