When your CAPM is not reliable anymore, Simulation and Optimization Tools will get the job done

Guilherme Marques Calôba, UFRJ
Regis da Rocha Motta, UFRJ
Alfredo Julio Souza Prates, Petrobrás
Marcelo Marinho Simas, Petrobrás

Abstract

In this paper we compare portfolio optimization by two distinct methodologies. The first one is the usual Capital Asset Pricing Model (CAPM) efficiency frontier, which considers that returns from every asset are normally distributed. Many apply this type of optimization to assets whose return distribution is simply non-normal.

We study two different cases: in the first one the distribution of the assets is normal and in the second one it is radically different. The point we are trying to make is that applying the standard CAPM to the non-normal distributions can lead to extremely erroneous decision-making. Furthermore, using software tools that combine optimization and simulation techniques, it is easy to optimize any portfolio, as long as its assets return distributions are known.

Conclusions and recommendations close the paper.

1 - Introduction

In a wide array of industries, the Net Present Value or Internal Rate of Return of a project turns out to be, considering uncertainties, approximately, a normal distribution. Some random disturbance on sales, costs (fixed and variable) and operations of sum and subtraction in order to achieve the annual cash flows end up in a normal distribution. This can be true for the so-called Opex-driven business, where investments, profits and costs have the same magnitude. In CAPEX-driven businesses, however, this is rarely the truth.

Oil industry, in particular, presents a very peculiar distribution pattern, generally following the oil-in-place distribution. As mentioned in Newendorp, this distribution tends to be a lognormal one. A brief explanation may come in handy.

The oil in place can be described in a very simplistic way as the following equation:

$OIP = A.h.\phi.(1 - Sw).B_{01}$

Where

OIP is the total volume of oil in place;

A is the rough area of the straight section of the reservoir

h is the height of the reservoir;

ϕ is the porosity factor, that is, the fraction of the rock that can hold oil;

Sw is the water saturation factor;

B_{01} is a volume formation factor, describing the volume variations due to temperature and pressure.

The recoverable oil is the product of the OIP and the recuperation factor. Figure 1 illustrates these considerations.

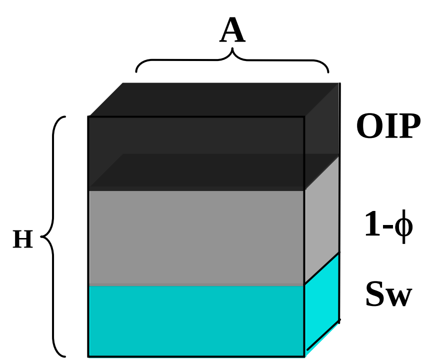

Figure 1 – Rough representation of oil in place (OIP) in a reservoir

As the oil in place is a product of a series of random variables, its logarithm is the sum of other variables (logarithms of the initial variables). This property points out to a lognormal distribution being a good proxy to the OIP distribution.

The problem poised in our hands is to optimize a portfolio of upstream opportunities considering their singular properties.

2 – Classic CAPM Portfolio Optimization

In dealing with normal distributed assets, the very well known Capital Asset Pricing Model developed by Markowitz in the early 50's. Considering a portfolio composed by two assets, namely 1 and 2, its return (μ) and risk (σ) can be calculated according to the following formulae:

$\mu_{Pfolio} = S_A.\mu_A + S_B.\mu_B$

$\sigma Pfolio = (S_A\sigma_A)^2 + (S_B\sigma_B)^2 + 2(S_{A.}\sigma_A).(S_B\sigma_B).\rho_{A,B}$

Where

Si is the share of asset I in the portfolio;

μ_i is the return (mean of net present value of the project, for instance) of asset i;

σ_i is the risk (standard deviation of the return) of asset i;

$\rho_{A,B}$ is the correlation between assets A and B.

This structure can be easily manipulated to obtain the so-called bullet-shaped curve, which is the efficient frontier of portfolio compositions, ranging from the smallest risk to the bigger return. This procedure can be done straightforward with any classical spreadsheet optimization tool, such as Excel Solver. The frontier is a non-dominated region, that is, one position in it cannot be considered superior to another unless some decision-maker utility is incorporated (as seen in Motta e Calôba).

In our example we consider four assets A, B, C and D. The formulae for return and risk can be thus generalized for this case:

$$\mu_{Pfolio} = S_A.\mu_A + S_B.\mu_B + S_C.\mu_C + S_D.\mu_D$$

$$\sigma_{PFolio} = \sqrt{\sum_{i=A}^{D}\sum_{j=A}^{D} S_i S_j .\sigma_i .\sigma_j .\rho_{i,j}}$$

The correlation between an asset and itself is a perfect positive one, so we always have $\rho_{i,i}$ equal to 1.

Considering this assets being normally distributed, the distribution for the portfolio return itself will also be a normal one, no matter the combination of the assets. In a more general case, this condition is not granted.

Our 4 projects, A, B, C and D are introduced in table 2, with information regarding its investment cost, its return measured as the net present value discounted for a minimum rate of atrativity and the risk as the standard deviation of the return of each asset:

Assets	Return	Risk	Investment
A	25,000	4,000	100,000
B	50,000	15,000	100,000
C	100,000	35,000	100,000
D	500,000	100,000	100,000

Table 1 – Summary of Investments Allowable – Values in US$ MM

We consider that our assets are statistically independent, that is, $\rho_{i,j}$ equals zero if j and I are different. We also introduce a budget constrain, allowing the investor to use only 200,000 million dollars.

Our procedure to estimating the frontier is a very simple one. First we minimize portfolio risk; then we maximize return, always subject to the constrains. We consider eight intermediate steps along the minimum risk and the one calculated when maximizing return, and maximize return in each of these nine problems.

Through non-linear optimization procedures internal to Excel Solver we obtained the following results, listing the return and risk along the frontier, as well as the share of the assets in each position. These results are shown in table 2.

	Portfolio Results		Portfolio Asset Composition			
	Return (μ)	Risk (σ)	A	B	C	D
Lowest Risk	91.007,97	14.231,67	100%	83%	15%	2%
2	147.562,64	18.853,36	100%	64%	23%	14%
3	204.117,31	28.370,92	82%	62%	32%	24%
4	260.671,98	38.547,01	43%	81%	43%	33%
5	317.226,65	48.844,21	4%	99%	55%	42%
6	373.781,32	59.502,48	0%	84%	63%	54%
7	430.335,99	70.716,36	0%	64%	70%	66%
8	486.890,66	82.261,06	0%	45%	78%	77%
9	543.445,33	94.014,80	0%	26%	85%	89%
Highest Return	600.000,00	105.948,10	0%	0%	100%	100%

Table 2 – Portfolios in the Efficient Frontier

As observed, higher risks can only be achieved by admitting higher return. As we rise in the frontier, the investor sweeps assets A and B, more conservative, for assets C and D, with higher return and risk. The efficient frontier can be observed in figure 2, stating the trade-off between risk and return.

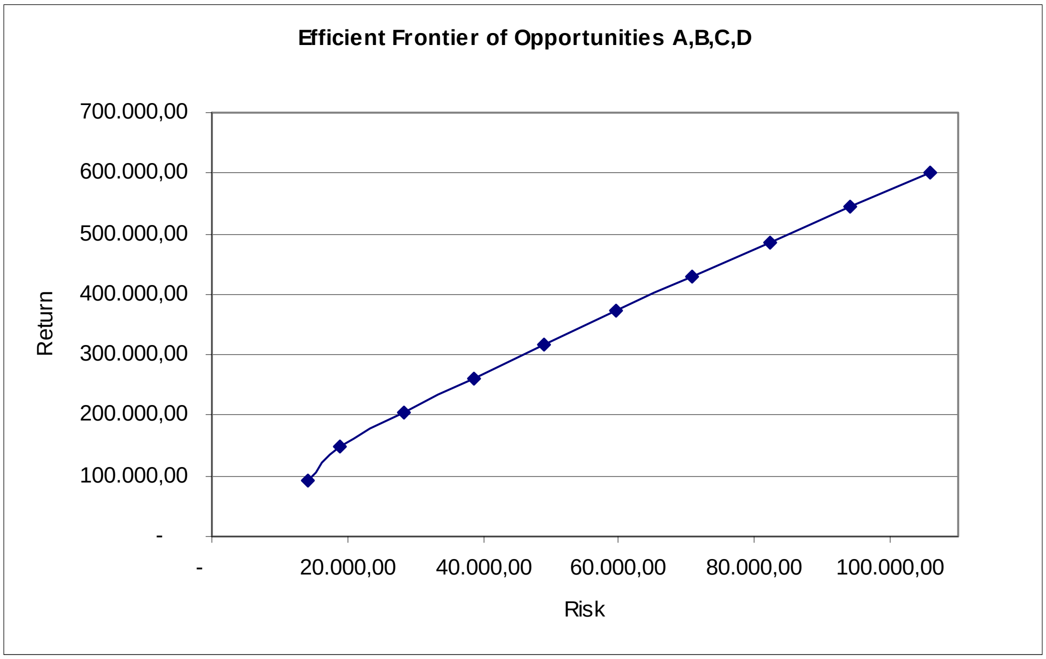

Figure 2 –Representation of the non-dominated portfolio compositions (Frontier)

If our assets are normally distributed, no longer discussion should be made. As this is not our case, we will move a little further, in the next session.

3 – Simulating our Portfolios

Considering our assets are lognormaly distributed in theory, we know we should not rely on the CAPM estimates. A little exercise will be done. Using @Risk, a Monte Carlo Simulation Add-in from Palisade Inc, we simulated the minimum risk portfolio pointed out by CAPM in two instances: the first one considering normal distributed returns and the second one considering the more realistic lognormal distribution. This will give us an insight in the importance of optimizing the portfolio considering appropriate distributions. We will consider the Value at Risk at a level of 5%, that is, the minimum profitability the portfolio yields with a 95% assurance. But first we shall show a graph of these two distributions, in figure 3.

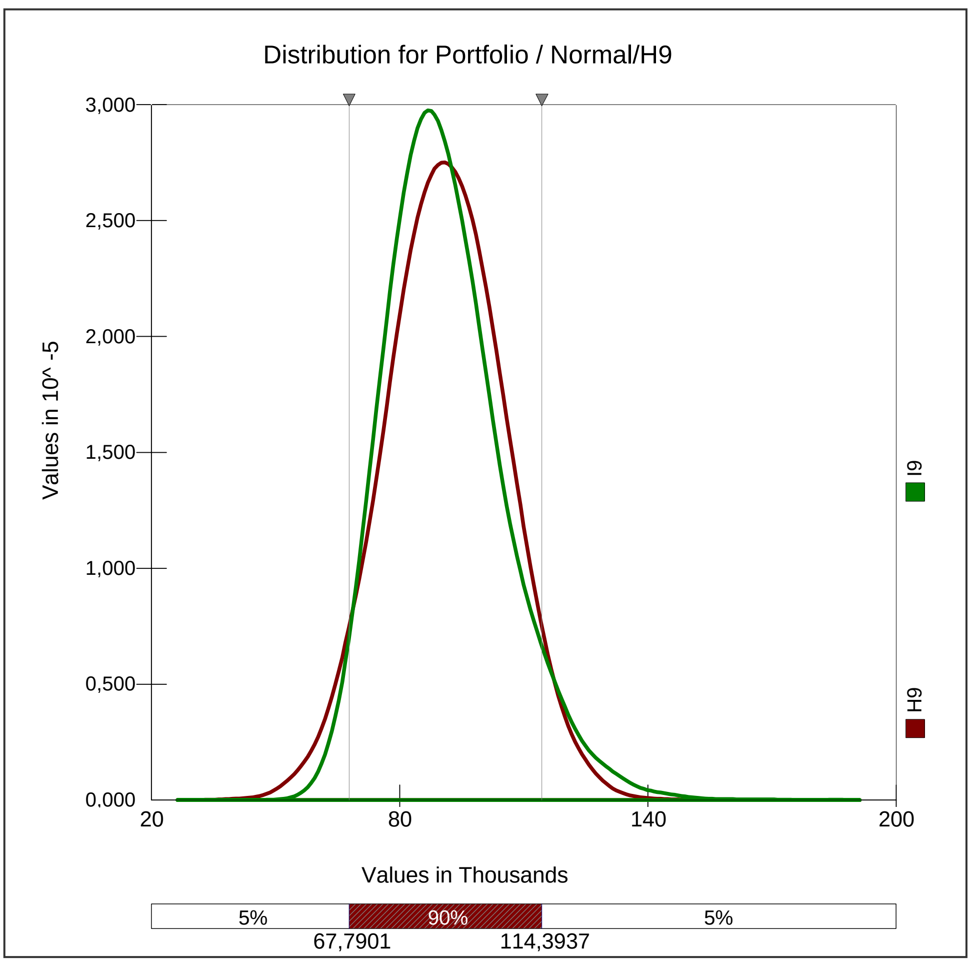

Figure 3 – A Normal and Lognormal Portfolio Distributions

As we can see, the normal distribution, in red pattern, has a higher mode, although it extends to lowers maximum and minimum. A list of descriptive statistics is shown in table 3, below.

Statistics	Normally Distributed	Lognormally Distributed
Minimum	25,881.34	51,060.59
Maximum	153,784.50	191,613.30
Mode	66,689.19	61,215.93
Average	91,090.78	91,034.36
Standard Deviation	14,212.45	14,333.36
Skewness	0.002369	0.6757061
Kurtosis	3.015074	3.894074
5% Percentile	67,790.07	70,322.95
95% Percentile	114,393.70	116,954.60

Table 3 – Statistics from Monte Carlo Simulation

We can observe a difference of approximately 2 billion dollars in both the 5% and 95% percentiles of the distribution. Higher differences appear in the mode, minimum and maximum values of the portfolios. However the investments have almost the same risk and return.

4 – Adapting for the Lognormal (or any other) Distribution

In the previous section we simulated the returns awarded by a portfolio composition that minimizes the risk. We noticed a difference in the portfolio return distribution functions when we swapped the distribution of the assets between normal and lognormal. However, using CAPM, we minimized risk for normally distributed assets. We wonder whether the composition of the risk-minimizing portfolio would change if we consider the assets lognormally distributed in our optimizing problem.

In order to investigate this issue, we would either have to conceive a way to represent theoretically the lognormal combination of the four assets composing the portfolio in order to optimize via Excel Solver or produce a means to optimize the portfolio while simulating. This last option seems very profitable as we could optimize portfolios composed by any kind of different distributions.

We couldn't figure out how to work this until we discovered the potential of Risk Optimizer, a component of Palisade's Decision Tools Industrial Version. This invaluable software can tackle this problem very easily. In fact, we can set the problem as easily as with Solver, determining our problem constrains, that is, budget allocation, and order to maximize or minimize a *statistic* of a given cell. If this cell contains an @Risk function, or is influenced by any random variables in the spreadsheet, the optimizer will do the trick.

Risk Optimizer works with state-of-the-art Genetic Algorithms (GA). This technique, in a nutshell, mimics Darwinian Evolution processes, creating subsequent generations of a population, which combines individuals in order to improve the desired result. The best individual is always kept in the next population. When the computer reaches a condition given by the user, such as running time or absence of improvement in the last x generations, the process is halted and we can see a log of the optimization and observe the values for the decision variables that provided this best result, as well as the whole optimization process. In order to work well, GA needs a feasible starting point in order to begin optimizing. Most of the cases this is no hard task to complete.

The introduction of this software creates a new paradigm in the spreadsheet technologies, that is, to easily optimize anything in any kind of spreadsheet with random variables. The GA optimization may take a long time, but this problem seems rather irrelevant with the processor speed in the computers nowadays. In addition to that, Palisade itself gave some attention to speeding up the simulation process, developing RiskAccelerator. There seems to be no limit for this tool.

In our case we used the same spreadsheet where CAPM was applied and defined the assets as lognormal distributions, with the aid of @Risk built-in functions that operate in perfect synchrony with Excel. The objective of the problem is to maximize return or minimize risk given the budgetary constrains. In a way very much similar to the approach we used in the CAPM optimization, we added one constrain commanding the return to stay above a given level and at the same time minimizing the risk for the portfolio. Rising this level from the one observed in the minimum risk up to the one of highest return, we obtain the efficient frontier.

Figure 4 shows the efficient frontier for the opportunities, considering the lognormal distribution of the opportunities.

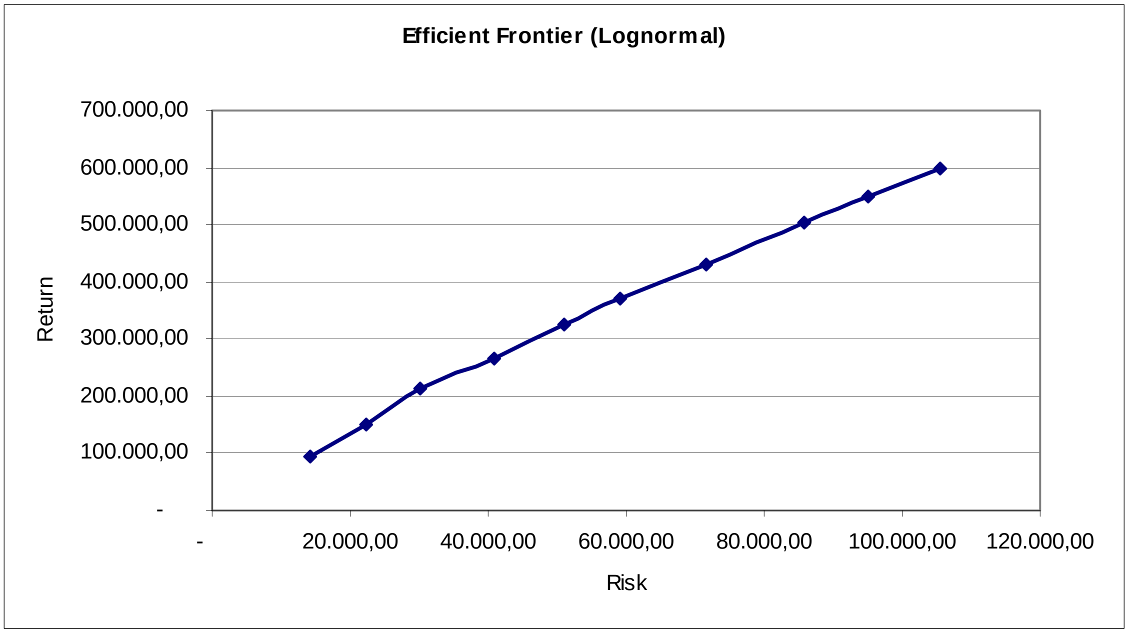

Figure 4 – Efficient Frontier observed via Simulation Optimization Techniques

We note that this frontier has some discontinuities, and the values observed for risk and return are not so different from the observed in the CAPM frontier. This can be easily seen when we plot both frontiers, which is done in figure 5.

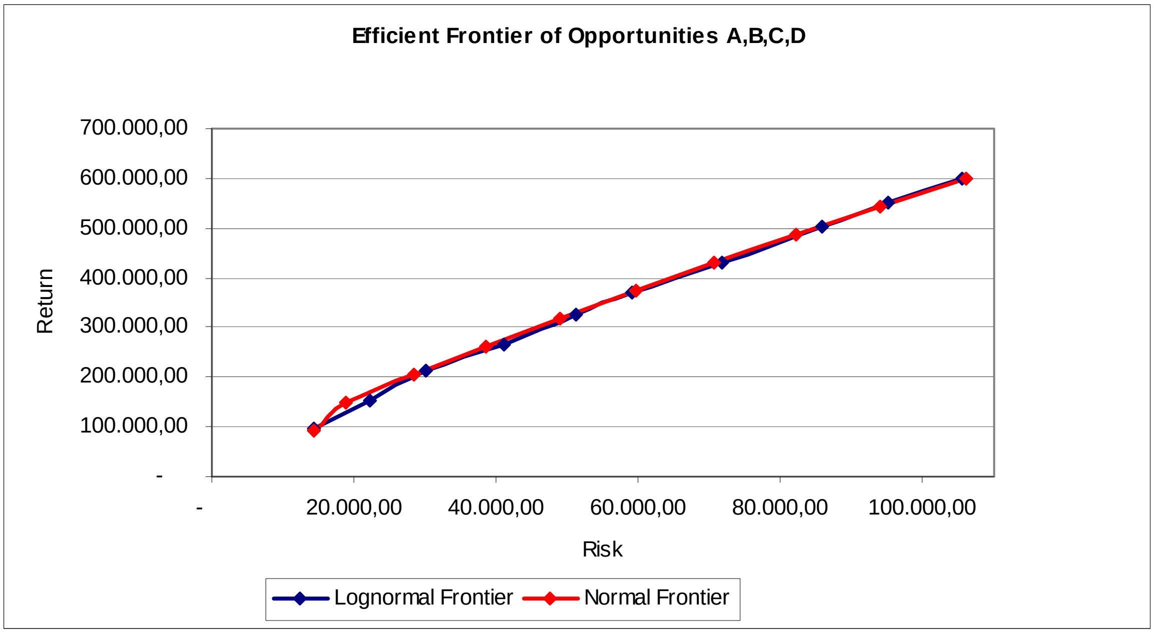

Figure 5 – Comparison of the Frontiers obtained via CAPM and Simulation

We may argue ourselves: we developed a numerical procedure in order to correctly optimize portfolios of 4 upstream assets, respecting the probability distribution of the assets NPV's, and for what? To reach the same efficient frontier? All is useless then. This *could* be our conclusion, but it would be a premature one.

And why is that? Maybe the frontiers are similar ones, but what about the portfolio compositions? Are they quite the same? That is, the same composition that minimizes risk / maximizes return for the normal assets does it again for the lognormal assets? Table 4 shows the composition of portfolios along the frontier obtained via simulation.

	Portfolio Results		**Portfolio Asset Composition**			
	Return (μ)	**Risk (σ)**	**A**	**B**	**C**	**D**
Lowest Risk	95,440.97	14,134.02	99,98%	81,36%	15,79%	2,88%
2	151,037.41	22,225.67	56,59%	94,50%	38,72%	10,19%
3	210,126.01	29,154.54	56,59%	90,23%	28,21%	24,97%
4	265,765.12	40,563.70	27,65%	70,19%	71,77%	30,40%
5	325,591.36	49,521.72	24,41%	90,80%	37,24%	47,54%
6	370,054.11	59,240.41	21,81%	69,51%	53,37%	55,31%
7	430,178.80	70,969.01	13,79%	74,69%	42,05%	69,47%
8	501,502.91	84,194.75	12,81%	23,59%	82,58%	81,02%
9	550,423.62	95,197.83	2,26%	11,96%	96,33%	89,46%
Highest Return	598,462.69	107,501.52	2,13%	1,55%	96,33%	100,00%

Table 4 – Portfolios in the Efficient Frontier (Simulation)

As seen before, as we rise in the frontier seeking higher returns, assets A and B positions are traded for positions in assets C and D. However the composition of the portfolios differ a lot.

Figure 5 shows a comparison between portfolio compositions along both for the normal and lognormal frontiers.

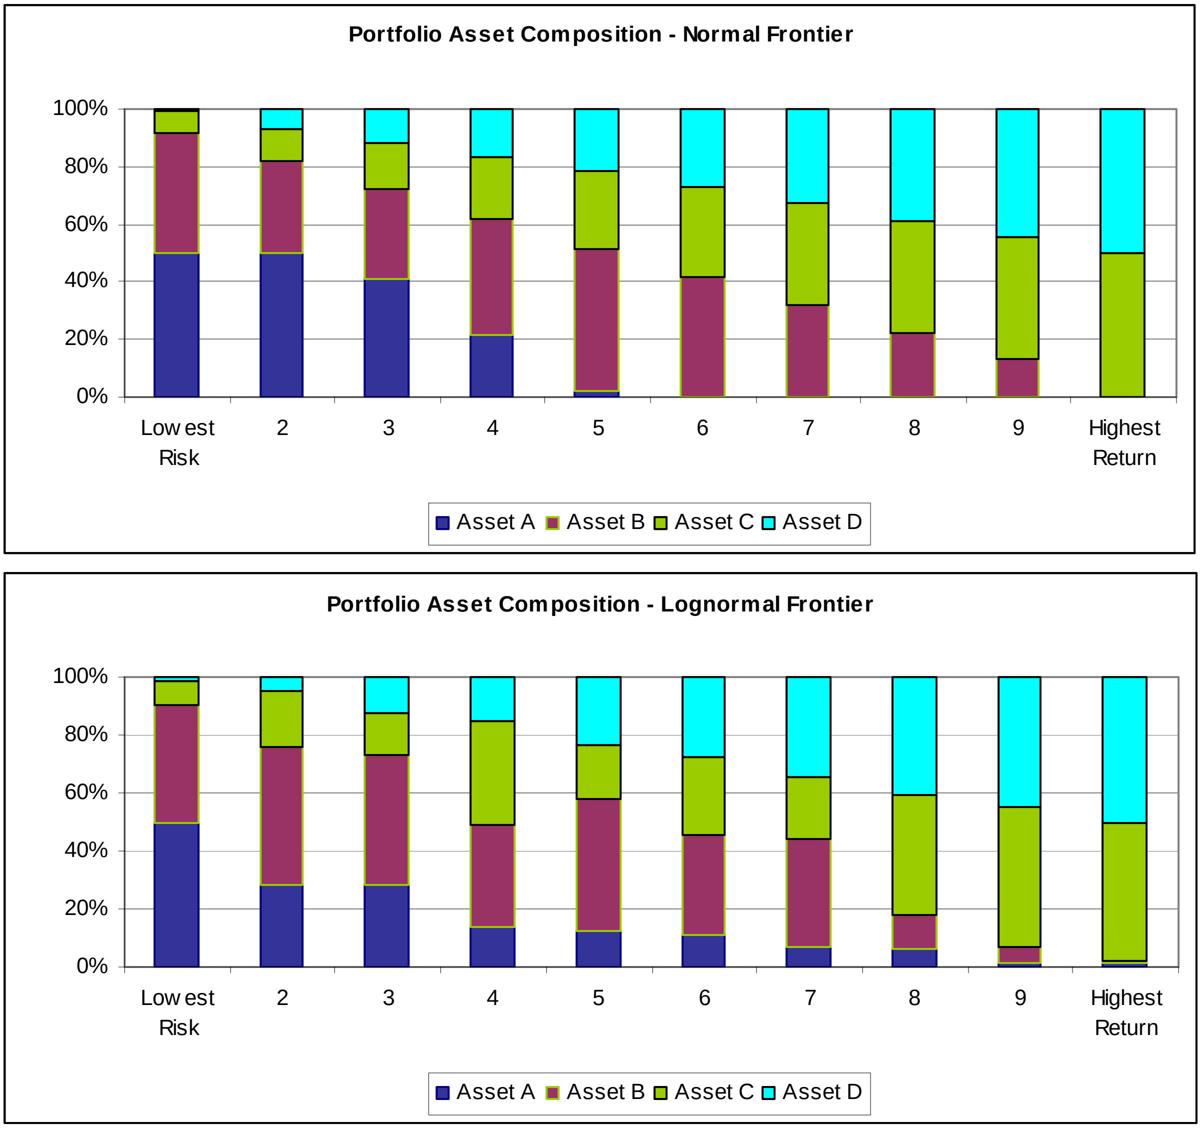

Figure 5 – Portfolio Composition along both frontiers

Observing figure 5 we can readily understand that, although risk and return are similar, the composition of the portfolio can be very different, as observed for points labeled 2, 4, 5, 6, 7 and 8. The figure confirms the existence of several differences between the two frontiers.

6 – A Final Comparison between the two techniques

As a finishing touch for this article, we shall compare the results obtained by both frontiers, that is, comparing graphs and statistics of the normal and lognormal frontiers. In a first step we will consider the minimum risk position in both frontiers, followed by the opposite point in the frontier, the highest return. This should provide us the information of the gains obtained from the use of the simulation-optimization method.

We shall emphasize that the aforementioned simulation method results are closer to reality as the assets introduced are considered lognormally distributed. This analysis will give us the precise dimension of the bias introduced by the CAPM. In order to do so, we simulated a 50,000-iteration run of the portfolio compositions.

As our first exercise, we introduce the graph obtained when simulating the minimum risk positions in both portfolios, labeled figure 6.

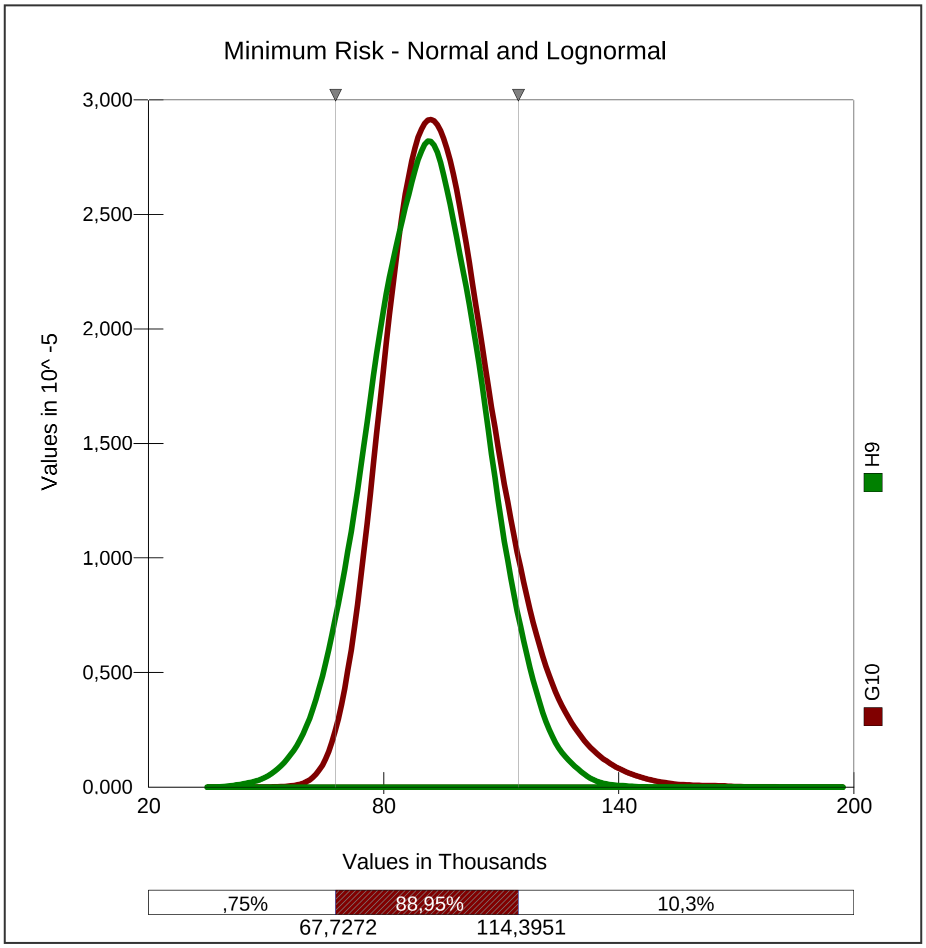

Figure 6 – Comparison between minimum risk positions

In the figure, the red pattern represents the lognormal distribution, and the green pattern corresponds to the normal distribution. We notice that the probability that the portfolio return is under 67,7 US$ MM is 5% for the normal distribution, but only 0,75% for the lognormal portfolio. Conversely, the probability of the portfolio awarding more than 114,4 US$ MM is also 5% for the normal distribution, but 10,3% for the lognormal distribution. In simple English, the CAPM amplifies the downside risk and reduces the upside potential for the portfolio.

The statistics that confirm this hypothesis are in table 5.

Statistics	Normally Distributed	Lognormally Distributed
Minimum	34.590,66	54.766,43
Mean	91.080,63	95.913,41
Mode	86.140,41	79.487,08
Maximum	148.052,70	197.638,60
Standard Deviation	14.183,18	14.399,29
Skewness	0,00	0,68
Kurtosis	2,97	3,94
5% Percentile	67.727,63	75.087,93
95% Percentile	114.392,30	121.710,80

Table 5 – Statistics from Monte Carlo Simulation

It is noticeable that the minimum value for the lognormal distribution is 20 million superior to the normal minimum. The mean of the lognormal portfolio is 4 million higher to the normal mean and the risk is more or less the same. The 5% percentile and the 95% percentile are also higher. We conclude that looking at the minimum risk portfolio through the CAPM glasses only worsens its value.

Moving up to the other side of the frontier, we simulate the highest return portfolio. The results are shown in figure 7.

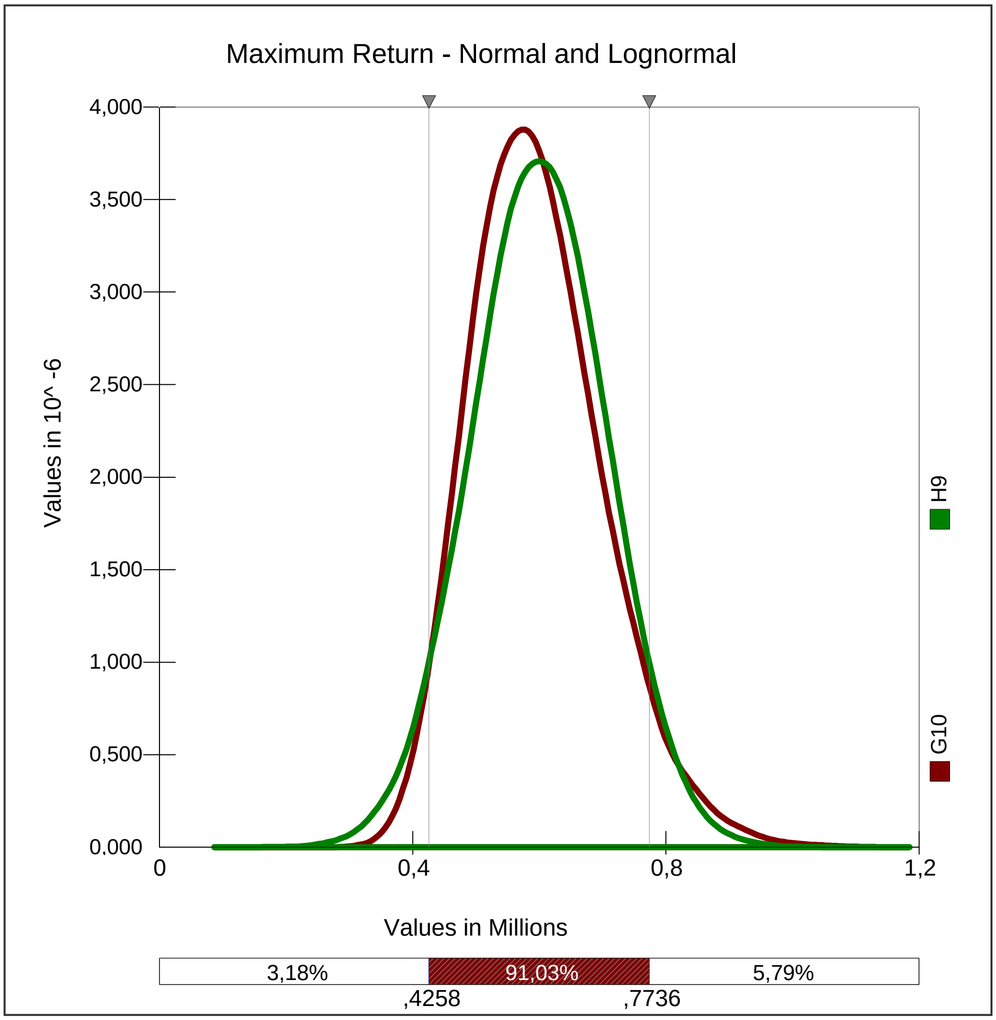

Figure 7 – Comparison between maximum return positions

The difference between the distributions is noticeable. Values under US$ 425 MM occur 3,18% (5%) in the lognormal (normal) portfolio. Values above US$ 773 MM appear 5,79% (5%) in the lognormal (normal) portfolio. Although the magnitude of the difference is smaller, the CAPM also reduces the value of this portfolio.

Table 6 shows statistics about those two portfolios:

Statistics	Normally Distributed	Lognormally Distributed
Minimum	83.489,61	286.708,10
Mean	600.071,50	597.300,90
Mode	545.298,20	548.532,80
Maximum	1.063.905,00	1.187.475,00
Standard Deviation	105.856,30	105.038,50
Skewness	0,00	0,54
Kurtosis	3,00	3,49
5% Percentile	425.701,30	441.457,40
95% Percentile	773.528,60	783.069,10

Table 6 – Statistics from Monte Carlo Simulation

Notice the huge difference between the minima of both simulations; it is worth US$ 200 MM! Differences in percentiles 5% and 95% are over US$ 10 MM. We can easily see that yes, there is a big difference between using CAPM and Simulation-Optimization Techniques and this difference can lead to improper selection of the portfolio.

7 – Conclusions

In this paper we introduced the necessity of using Simulation-Optimization techniques to handle portfolio management in the absence of the heavenly desired normal distributed assets.

We introduced four assets slightly based on Upstream Exploration and Production Oil Projects, assuming lognormal distributions. Throughout the paper, using Solver to compose the CAPM efficient frontier of the assets and Palisade’s RiskOptimizer in order to compose the Simulation-Optimization frontier of these same assets. We’ve shown that composition of portfolio may vary a lot in both cases.

Lastly, we simulated both portfolios, using Palisade’s @Risk, and showed the relevant differences between the theoretical results and the Simulation-Optimization Techniques.

This academic example has another point: stimulating people to use RiskOptimizer and discover how the easy-to-use, practical approach of Genetic Algorithms Optimization comes as a interesting solution not only for portfolio or financial problems, but any problem at all involving optimization of a spreadsheet containing random variables, for the generic approach of this great software may be inserted in a huge number of problems.