

Massachusetts Institute of Technology

ESD.71 – Engineering Systems Analysis for Design

Application Portfolio

Heating System for a Group of Condominiums

Rory Clune

December 2008

Executive Summary

In the context of a natural gas-fired heating system for a group of condominiums in the West of Ireland, the impact of uncertainty and the value of incorporating flexibility in design are investigated. The heating system is modelled using a combination of real input data and typical performance parameters reflective of such a system. The variability in demand for heat, taken to be directly proportional to the number of tourists visiting the region, and the price of natural gas are the key uncertainties considered.

The net present (monetary) values of the system with and without flexibility in design are evaluated using a two-stage decision tree over ten years and a binomial lattice over six years. In both cases, the incorporation of flexibility increases the expected NPV of the system. Other decision-making criteria besides NPV are explored. The results of these analyses are presented and discussed.

Contents:

1. Introduction and problem outline
2. Recognition of uncertainties
3. Formal system definition
4. Decision analysis
5. Binomial lattice analysis of evolution of a major uncertainty
6. Binomial lattice valuation
7. Conclusion
8. Appendices

1. Introduction and problem outline

A new building technology company, *HeatWise Plc*, has been awarded the contract to provide heat to a group of fifty condominiums built just outside Killarney in Ireland, an area noted for attracting national and international tourism. The condominiums are to be rented to tourists visiting the area, and will be marketed as a more ‘authentic’ Irish experience than a standard hotel. Given the country’s temperate climate, most buildings in Ireland are heated almost year-round – especially buildings accommodating tourists who are used to warmer climates. Air conditioning is rarely employed in residential settings, and it is only in the height of summer that buildings are not heated at night.

The conditions of the contract stipulate that the company will provide a service which has a certain minimum heating capacity, as defined by the developer of the condominiums. *HeatWise* sells the heat by the kilowatt-hour to the consumer – the developer has little experience of building technology and does not wish to manage such a system. In the case of the heating capacity of *HeatWise*’s plant being exceeded by the heating demand, the developer has arranged for an outside source to bring in separate storage heaters. In such a scenario, *HeatWise* would miss out on additional revenues, but the buildings would not go unheated.

The developer has indicated that the current group of fifty condominiums represents only the first phase of construction, and that he intends to construct another phase five years after the first. The size of this second phase of condominiums will be determined by his perceptions of the demand for the first phase and by his predictions of future tourism levels. It is, therefore, entirely possible that none of this second phase will be built. It has been agreed that the heating contract for the second phase will initially be offered to *HeatWise* before being offered to the market. The total amount of condominiums he can build is constrained by planning regulations to eighty. For the sake of simplicity, the second construction phase is considered to add either thirty, ten or zero condominiums to the scheme. The rationale for this is explained in a later section.

1.1 The heating system - What does it include, and what does it exclude?

The system in this project is defined as the natural gas boiler plant which *HeatWise* has selected for installation. The plant will heat water in a gas boiler system, and pipe it through a closed network to the condominiums.

The revenues generated are proportional to the heat delivered to each condominium unit, as determined by a fixed cost per kilowatt-hour charged to the consumer. The cost of generating the heat is dependent on the cost of natural gas. What makes this problem interesting is that the efficiency (simply defined as heat output divided by heat input) of a gas boiler declines as the load on the boiler decreases below the design load. This adds an extra degree of complexity to the system, as the relationship between the cost and the amount of heat supplied is nonlinear.

1.2 Principal design levers

Two distinct design scenarios are envisaged:

- (1) The construction of a plant comprising a single large boiler, capable of meeting the heating demand from all possible eighty condominiums at winter design temperatures.

- (2) The construction of a plant building of similar geometric proportions, but with a smaller boiler capable of meeting the heating demand for the first phase of construction only. Should the developer go ahead with subsequent phases, the system can be expanded by installing a second boiler into the network, capable of meeting the demand from 30 additional units.

It is assumed that the initially installed piping network to the boiler plant room remains unchanged after initial construction, as the retrofitting of such a network upon plant expansion would be entirely impractical.

1.3 Numerical model

Based on the author’s previous experience of building technology projects, a model is constructed in Microsoft Excel. The model determines the heating load of a typical condominium based on monthly average temperatures for the region. This can be converted to a set of requirements for the gas boiler.

The model as described so far will, for all practical purposes, act as a ‘black box’ which relates the demand on the system (and, of course, the revenue generated) to the total cost of generating the heat to meet this demand, comprising an initial capital investment plus a stream of annual/monthly variable costs over the system’s design life. Further details of this model are explored at a later stage.

The benefit of the system to *HeatWise* is the financial profit generated by its operation, defined here as the net present value of costs and revenues. It is not intuitively clear at the outset how the compounding factor of variable boiler efficiency will affect the end results.

Year	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Number of condominiums	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50
Number of condos built	50														
Total Heat Load (kW)		355	360	335	303	253	193	155	155	188	238	308	335	355	360
Boiler Capacity (kW)	360	360	360	360	360	360	360	360	360	360	360	360	360	360	360
Capacity Added (kW)	360														
Plantroom size (kW equivalent)	576														
Part Load Ratio		0.99	1.00	0.93	0.84	0.70	0.53	0.43	0.43	0.52	0.66	0.85	0.93	0.99	1.00
Heat Input (kW)		398	400	392	382	368	353	345	345	352	364	383	392	398	400
Efficiency		89%	90%	85%	79%	69%	55%	45%	45%	53%	65%	80%	85%	89%	90%
Number of hours per month		250	220	190	160	130	100	70	40	100	160	220	280	250	220
Fuel In (kWh)	0	104,827	92,632	78,372	64,312	50,321	37,140	25,394	14,511	37,019	61,271	88,776	115,495	104,827	92,632
Cost of gas (\$/kWh)		0.066	0.067	0.067	0.068	0.068	0.069	0.069	0.070	0.070	0.071	0.071	0.072	0.073	0.074
Capital Cost (\$)	41,910	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Variable Cost (\$)	0	6,976	6,214	5,299	4,383	3,457	2,572	1,773	1,021	2,626	4,381	6,398	8,390	7,676	6,837
Discounted Costs (\$)	41,910	6,918	6,111	5,169	4,240	3,317	2,447	1,673	955	2,437	4,032	5,840	7,595	6,891	6,087
NPV of costs (\$)	540,298														
Revenue (\$)	0	26,625	23,760	19,095	14,520	9,848	5,775	3,255	1,860	5,625	11,400	20,295	28,140	26,625	23,760
Discounted Revenues (\$)	0	26,405	23,369	18,625	14,046	9,447	5,494	3,071	1,741	5,220	10,492	18,524	25,473	23,902	21,154
NPV of revenues (\$)	1,077,126														
NPV (\$)	536,828														

Fig 1.1 Spreadsheet model screenshot

1.4 Contextual factors that will affect the value of system performance

- (1) The demand for the condominiums and, by extension, for heat, is crucial in affecting the system's performance. Data on Irish tourism is readily available through government agencies such as the Central Statistics Office of Ireland. The area in which the condominiums are to be built is noted for the number of visitors it attracts from the US. Tourism in Ireland is affected by a plethora of factors, among them the cost of living in the country and worldwide economic cycles.
- (2) The future price of natural gas, the fuel source for the central plant, is the second key contextual factor which will affect the system's performance. The heating contract makes no provision for the increase of prices as the cost of gas rises.

It is recognised that a number of other uncertainties may come into play here. Annual variations in average outdoor temperatures, which would affect the number of degree days (a measure of the amount of heat required throughout the heating season) are likely to play a role. This source of uncertainty, and a number of others considered, were felt to be less crucial to the system than the two chosen above. As such, their effects were neglected

2. Recognition of uncertainties

The two primary uncertainties which exist over the lifetime of the system are now formally defined.

2.1 The demand for the heat from the system

As previously discussed, the demand for the system is directly related to Irish tourism levels. The model developed allows for units of demand to be added or removed, in the form of discrete individual condominiums. Based on the heating demand of a typical condominium, this then translates into a tangible kW and kWh (kilowatt-hour) demand on the heating system.

Information on Irish tourism is widely available. Historic data of number of visitors and the amount of time and money they spent are available from Ireland's Central Statistics Office's (CSO). Annual data on the number of people entering the country for tourism purposes since 1995 are sourced from the CSO website¹. Using logarithmic regression methods, the following parameters are extracted from the data:

-mean growth rate per annum = 3.63%

-volatility = 6.74%

The exponential curve fitted to these data can then be used to predict future levels of tourism, and the volatility provides upper and lower bounds.

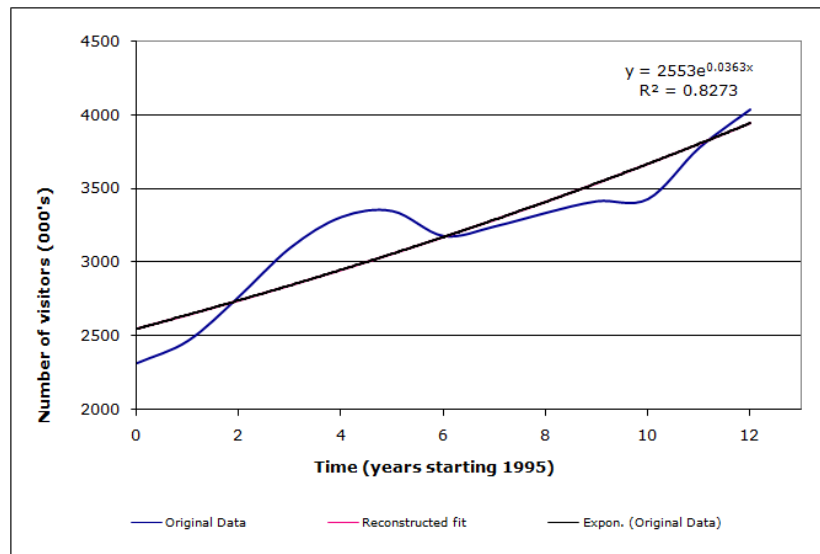


Fig 2.1 Regression analysis of number of tourists

In relation to the heating system, it is assumed that the initial number of condominiums built will be exactly enough to accommodate that level of tourism demand. The relationship between the number of tourists visiting the country (available data) and the number of condominiums built is linear. i.e. if 1000 more people visit the country, 1 more condominium will be built etc. Thus, the demand for the

¹ <http://www.cso.ie/px/pxeirestat/database/eirestat/Tourism/Tourism.asp>

condominiums (and, by extension, heat for the condominiums) is reasonably assumed to also have an annual mean growth rate of 3.63%, and a volatility of 6.74%

2.2 The price of the natural gas that fuels the system

The price of natural gas is taken as the price for residential natural gas listed by the Energy Information Administration². Monthly data for the last two years are taken, and the same regression procedure as before is applied. The regression shows a mean monthly growth rate of 0.90%, with an associated volatility of 14.0%. This is reflective of a large uncertainty concerning the future price of natural gas, owing to the large variations in recent historic prices.

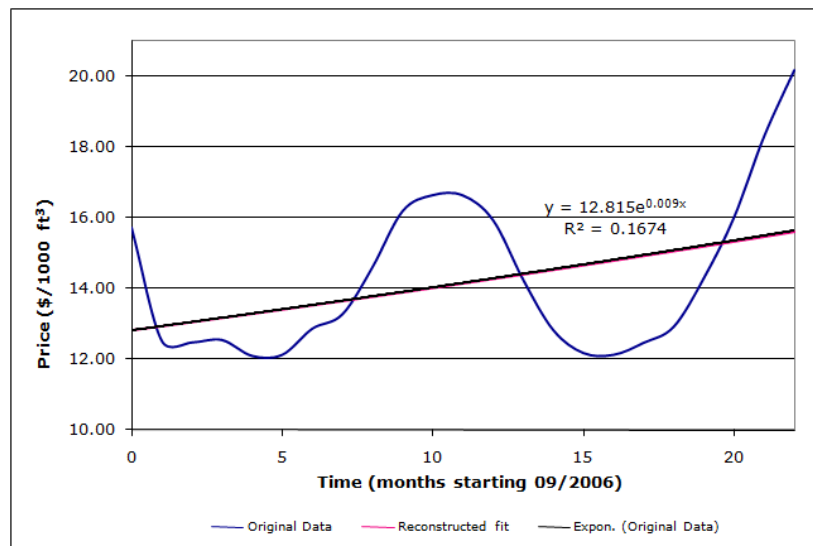


Fig 2.2 Regression analysis of price of natural gas

The generated exponential model of price is then used to determine an expected future price, and the volatility numbers are the upper and lower bounds on that expectation. The R^2 value (a measure of how well the reconstructed fit matches the original data) is particularly low, which suggests that an exponential model may not be entirely appropriate in this context. This is due to the large fluctuations in the original data which can be seen above. When the broad trend of prices over a number of years is studied², it is apparent that the underlying trend behind these large fluctuations can be represented exponentially. A high-order polynomial curve, for example, may capture the large fluctuations above better than an exponential curve, thus reducing the R^2 value. It was felt that such an approach would miss the underlying trend which will be extrapolated much further into the future during subsequent analyses.

² <http://tonto.eia.doe.gov/dnav/ng/hist/n3010us3m.htm>

3. Formal system definition

The following defines the two major system concepts employed in the course of this portfolio, both of which can be accommodated by the developed model. Before describing these, it will be instructive to briefly explain the initial capital cost function used in the model. It is of the form:

$$C(Q, P) = AQ^b + DP$$

Where Q is the output heat capacity of the installed boiler, P is the size of the plant room built to house the boiler (and any future boilers that may be installed), A and D are scaling factors dependent on the units of Q and P, and b is an economies of scale factor. The model is currently run using an economies of scale factor of 0.7 This cost function represents the capital costs of the system only. The variable costs are given by the (uncertain) price of natural gas times the amount of natural gas required to fuel the system. The revenues accruing to the system are given by a fixed charge multiplied by the amount of heat supplied to customers. Maintenance costs are not accounted for.

3.1 Fixed Design

It is assumed that, halfway through the ten-year design life of the system (at t=5), an additional ten condominiums will be built. The initially installed boiler is capable of meeting the heating load from this predicted future increase in demand. Considering this design qualitatively, it can be seen that it takes advantage of the economies of scale associated with purchasing the bigger boiler immediately. However, the efficiency of a gas boiler varies with the load placed on it³. Thus the boiler will be running at a relatively lower efficiency before the extra demand is placed on it. There are, of course, many possible variations of this design. The deterministic case outlined above can vary enormously, depending on the predicted timing and magnitude of any demand increments. Despite these variations, the key characteristic of this fixed design is that a single boiler is installed at the outset, with no provision for further increases in capacity.

3.2 Flexible Design

A smaller boiler, the capacity of which is designed in response to the current heating demand from fifty condominiums, is installed. It is envisaged that, if any extra condominiums are built, causing an increase in future demand, another boiler can be then installed. There is, however, a price to pay for this flexibility. The plant room (of size P) built at time t=0 must be large enough to accommodate the second boiler which may or may not be installed. This increased plant room cost also includes the cost of laying the pipes and other necessary features for the 'future' boiler. The cost model incorporates a second capital investment at t=5, recognising the cost of installing additional capacity.

This design does not take advantage of the economies of scale involved in purchasing a large boiler to the same extent as the fixed design. By contrast, as well as having the advantage of being able to respond to uncertain changes in demand, the system using a smaller boiler will initially operate at a higher efficiency than an oversized boiler. As in the fixed case, many variants of this design are possible. The key characteristic is that the initially installed boiler is of a capacity which is at least enough to meet the initial demand, while extra plant room space is constructed in anticipation of the purchase and installation of boilers in response to future changes in demand.

³ An appendix outlining the theory of boiler efficiency is provided

4. Decision Analysis

Explanation of decision tree structure

The decision tree, included in the accompanying Excel file and shown in fig 4.3, is structured in the following way:

At time $t=0$, a decision is made between the two major designs identified in section 3. To recap, the choice is between a fixed and a flexible design. The fixed design consists of installing a boiler and constructing a plant room capable of meeting the forecast future demand after five years. Initial demand comes from 50 building units, and the forecast demand growth case anticipates that another 10 units will be built after the initial five-year period has elapsed. The flexible design consists of a boiler which is capable of meeting only current (50 units) demand. Additional plant room space and infrastructure is incorporated to accommodate the possible future installation of an additional boiler, increasing the initial capital costs. It is interesting to note, however, that the capital cost for the fixed design case is almost equal to the capital cost for the flexible case – the cost of incorporating flexibility (plant room space, piping etc.) is offset by the fact that the flexible case starts out with a smaller, cheaper boiler than does the fixed case. Reductions in plant room costs are offset by differences in costs of the boilers.

The two uncertainties previously described now come into play. The uncertainty surrounding both the cost of natural gas and the future demand (tourism) is defined here by three cases – forecast, favourable and unfavourable. The exact levels and associated probabilities are derived from a normal probability distribution fitted from the results of the regression analysis performed in section 2. The monthly growth rate for the price of gas is either -0.90%, 0.90% or 2.00%, and the increment in demand after five years is either 0, 10 or 30 additional building units. This leads to a total of nine possible scenarios, and each of these is run through the model.

After five years, a second decision must be made. This decision relates to expansion of the system. The choices are to remain unchanged, to install a medium-sized boiler (capable of servicing an additional 10 units), or to install a large-sized boiler (capable of servicing an additional 30 units) In the case of the fixed design, no second-stage decision can be taken, and the system remains unchanged. After this decision, the same uncertainty criteria apply for the next five years, bringing us to the end of the system's 10-year lifespan. The analysis does not recognise the additional complexity of the uncertainty criteria changing for the second five-year period. Outcomes expressed in terms of NPV are determined for all possible outcomes.

Using the folding back technique of decision analysis, the individual second-stage decisions delivering the best expected outcomes are identified. A number of interesting points emerge from this analysis. The behaviour of the fixed design is as expected. It responds well to favourable conditions, and badly to unfavourable conditions. An expected value is calculated by multiplying each outcome by the probability of its occurrence and summing.

The flexible design results are more intriguing. In the case of forecast conditions occurring, the best decision is in fact to 'stay', rather than to install an additional boiler to meet the demand from 10 additional units built in is 'forecast' case. This can be understood by considering that the demand on the heating system varies seasonally. The high demand and associated profit in winter is outweighed by the inefficiency of a larger system running well below capacity during warm months, and the optimal policy turns out to be not to increase capacity. When demand conditions are favourable, and demand increases due to 30 extra buildings in year 5, the benefit of the extra revenues outweighs this

inefficiency, and the best policy becomes the increasing of capacity to fully match additional demand. As expected, in unfavourable conditions, the best decision is generally not to expand.

The VARG curve below shows the probabilistic cumulative distribution of NPVs for each design case. It is readily identified that the flexible design is better than the fixed design regardless of outcome, and that the benefit is especially pronounced at the lower end of the spectrum. Thus, the incorporation of flexibility design is seen to dramatically reduce the potential downside for the business, as well as to increase the expected NPV and potential upside.

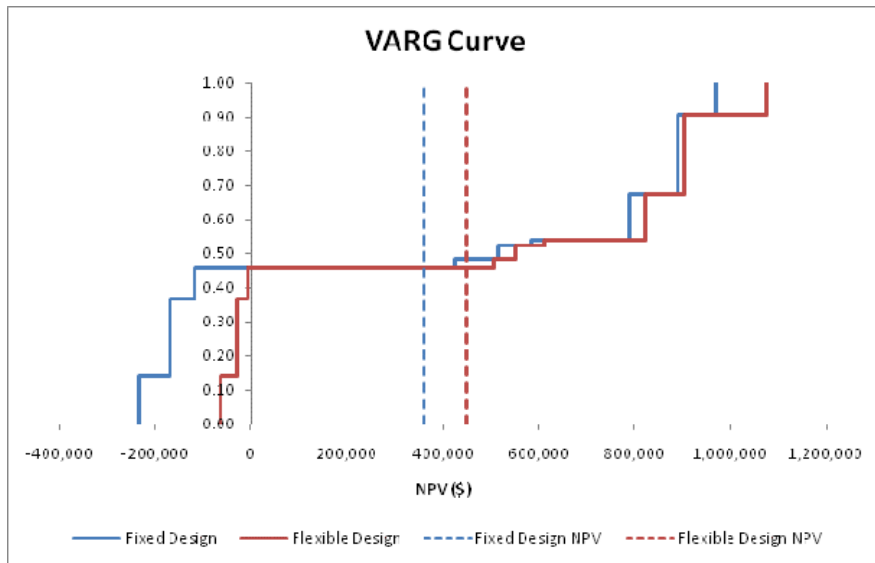


Fig 4.1 VARG curve from decision analysis

The following table identifies a number of decision criteria which may be used, all of which the system designer should be conscious. The flexible design is seen to perform better than the fixed one on all counts, with the exception that its initial capital expenditure is slightly larger (as discussed previously).

	Fixed Design	Flexible Design	Better?
ENPV	360,667	447,290	Flexible
NPV _{max}	968,957	1,075,188	Flexible
NPV _{min}	-235,821	-65,301	Flexible
CAPEX	40,943	41,910	Fixed

Fig 4.2 Multiple criteria for decision making

Rory Clune - AP4 -10/24/08		D1	Uncertainties (Chance nodes)					D2					
Expected Value (\$)		First Decision	Expected Value (\$)	Chance Node 1 (Gas Price)	Probability	Chance Node 2 (Tourism)	Probability	Outcome (\$)	Expand Decision	Outcome (\$)			
447,290	Start	Fixed Design	360,667	Forecast	0.08	Forecast	0.50	513,481	NA	513,481			
	Favourable					0.20	583,344	NA	583,344				
	Unfavourable					0.30	424,608	NA	424,608				
	Favourable			0.46	Forecast	0.50	Forecast	0.50	890,103	NA	890,103		
							Favourable	0.20	968,957	NA	968,957		
							Unfavourable	0.30	790,294	NA	790,294		
	Unfavourable			0.46	Forecast	0.50	Forecast	0.50	-169,221	NA	-169,221		
							Favourable	0.20	-118,044	NA	-118,044		
							Unfavourable	0.30	-235,821	NA	-235,821		
	Flexible Design			447,290	Forecast	0.08	Forecast	0.50	Forecast	0.50	551,235	Stay	551,235
									Favourable	0.20	610,549	Stay	539,364
									Unfavourable	0.30	506,011	Stay	432,008
		Favourable	0.46				Forecast	0.50	Forecast	0.50	905,160	Stay	583,117
									Favourable	0.20	1,075,188	Stay	609,227
									Unfavourable	0.30	822,386	Stay	610,549
		Unfavourable	0.46				Forecast	0.50	Forecast	0.50	-32,331	Stay	506,011
									Favourable	0.20	-8,895	Stay	450,491
									Unfavourable	0.30	-65,301	Stay	340,948
		Forecast	0.08				Forecast	0.50	Forecast	0.50	551,235	Go medium	873,503
									Favourable	0.20	610,549	Go medium	905,160
									Unfavourable	0.30	506,011	Go big	875,382
	Favourable	0.46	Forecast	0.50	Forecast	0.50	905,160	Go big	909,448				
					Favourable	0.20	1,075,188	Go medium	984,014				
					Unfavourable	0.30	822,386	Go big	1,075,188				
	Unfavourable	0.46	Forecast	0.50	Forecast	0.50	-32,331	Go big	822,386				
					Favourable	0.20	-8,895	Go medium	805,351				
					Unfavourable	0.30	-65,301	Go big	775,057				
	Forecast	0.08	Forecast	0.50	Forecast	0.50	551,235	Stay	-32,331				
					Favourable	0.20	610,549	Go medium	-133,044				
					Unfavourable	0.30	506,011	Go big	-399,282				
	Favourable	0.46	Forecast	0.50	Forecast	0.50	905,160	Stay	-8,895				
					Favourable	0.20	1,075,188	Go medium	-81,867				
					Unfavourable	0.30	822,386	Go big	-264,051				
	Unfavourable	0.46	Forecast	0.50	Forecast	0.50	-32,331	Stay	-65,301				
					Favourable	0.20	-8,895	Go medium	-199,644				
					Unfavourable	0.30	-65,301	Go big	-471,468				

Fig 4.3 Decision Tree

5. Binomial Lattice Analysis of Evolution of a major uncertainty

Price of natural gas

The uncertainty analysed in the lattice analysis is the price of natural gas – a key input to the system. The lattice analysis disregards any uncertainty surrounding the demand levels. In fact, since the lattice analysis considers only a five-year period, the simplifying assumption is made that no additional condominiums will be built in the second phase.

5.1 Modal forecast and volatility

The analysis begins with the most recently available price of \$20.19 / 1000 ft³, and uses the trend extracted from available the data in section 4. A mean monthly growth rate of $v_m = 1.17\%$ and an associated volatility of $\sigma_m = 13.86\%$ were extracted in that analysis. These are converted to an annual growth rate and associated volatility by the following relations:

$$v = (1 + v_m)^{12} - 1 = 14.98\%$$

$$\sigma = \text{sqrt}(12\sigma_m^2) = 48.01\%$$

The expected price trend over time is therefore of the form:

$$P(t) = 20.19e^{0.1498t}; \text{ with } t \text{ in years}$$

There is a volatility of 48.01% around this mean, and the Δt time step for the lattice analysis is one year.

5.2 Developed Lattices

The u,d and p parameters are derived from the above information as in the class example. The relevant relationships are:

$$u = e \exp(\sigma\sqrt{\Delta t})$$

$$d = e \exp(-\sigma\sqrt{\Delta t})$$

$$p = 0.5 + 0.5(v/\sigma)\sqrt{\Delta t}$$

The following two lattices (outcome & probability), representing the uncertainty evolution over a five-year period are developed:

Outcome Lattice					
20.19	32.63	52.74	85.25	137.78	222.70
	12.49	20.19	32.63	52.74	85.25
		7.73	12.49	20.19	32.63
			4.78	7.73	12.49
				2.96	4.78
					1.83

Fig 5.1 Lattice of outcomes (price per 1000 ft³ of gas)

Probability Lattice					
1.00	0.66	0.43	0.28	0.19	0.12
	0.34	0.45	0.44	0.39	0.32
		0.12	0.23	0.31	0.33
			0.04	0.11	0.18
				0.01	0.05
					0.00

Fig 5.2 Lattice of probabilities

The PDF of possible gas prices at year 5 is shown on the following plot:

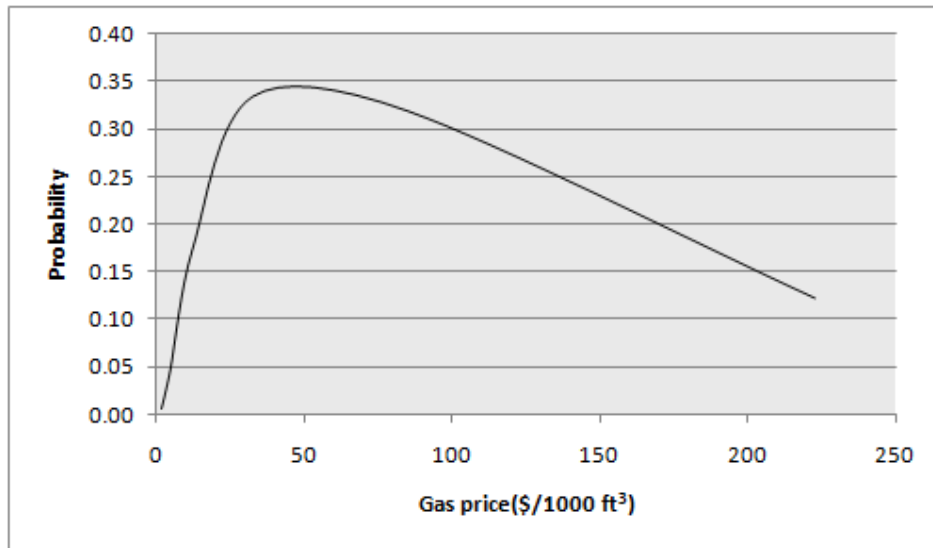


Fig 5.3 PDF of gas prices at year 5

6. Binomial Lattice Valuation

The lattice of possible natural gas prices developed in the previous section is now taken forward. Slight modifications are made to the numerical model in order to facilitate the lattice analysis techniques. A single year of monthly cash flows is analysed, with a mean monthly growth rate accounting for price changes throughout the year. The prices from the lattice of section 5 become the prices at the start of each year, and the model then accounts for the variations throughout the year. As mentioned before, the demand level remains constant from year to year, equal to the initial demand from the first phase of construction.

A lattice of cash flows is presented below – each entry represents the undiscounted net cash flow derived in a single year, given the various prices. The first entry in the lattice includes the initial capital investment of approximately \$40,000, while all subsequent entries include only the revenues and costs which occur in that particular period. The lattice incorporates six years of cash flows, since year 0 (the first entry in the lattice) also has a cash flow associated with it (this is a natural result of the way the numerical model is used to account for price increases throughout a year)

Lattice of Cashflows					
61,128	61,650	5,245	113,365	288,118	570,567
	128,646	103,039	61,650	5,245	113,365
		144,489	128,646	103,039	61,650
			154,292	144,489	128,646
				160,356	154,292
					164,109

For clarity of presentation, these figures are now expressed in thousands.

Lattice of Cashflows ('000s)					
61.1	61.7	5.2	113.4	288.1	570.6
	128.6	103.0	61.7	5.2	113.4
		144.5	128.6	103.0	61.7
			154.3	144.5	128.6
				160.4	154.3
					164.1

Fig 6.1 Cash flow lattices

6.1 Flexible Analysis

The dynamic programming (backtracking) procedure is now applied. There is a put option on the system, in the form of the possibility to close it down, immediately stopping all future costs and revenues. There is no ongoing penalty cost associated with a system that has been closed down. Once the system has been closed down, it cannot be opened again. This ensures path independence.

Each entry in the dynamic programming lattice is modified to represent the cash flow from that period plus the maximum possible expected value from all following periods, given the option to close. This recursive approach, working from the end to the beginning, gives the NPV of running the system over a six year period in the first cell.

As an example, the expected cash flow from year six seen from the top cell of year 5 in fig 6.2 (entry = \$288.1) is the maximum of (-\$570.6 times the probability of the price going up plus -\$113.4 times the probability of the price going down) and zero (the cash flow if the system is shut down). The NPV of this cash flow as seen from year 5 is then the cash flow discounted by one year at a discount rate of 10%.

$$E(NPV_{\text{year } 6}) = \max[(-570.6 \cdot 0.66 + -113.4 \cdot (1-0.66)), 0] \cdot (1+0.10)^{-1}$$

$$= 0 \cdot (1+0.10)^{-1} = 0$$

Therefore, the value associated with being in the top cell of the lattice column in year 5 is \$288.1 plus \$0, or just \$288.1 (remembering that these values are presented in thousands of dollars). This backtracking process is applied to all cells, looking one year in advance in each case. The cumulative nature of this process gives the NPV of the project accounting for the option in the first cell of the lattice, as shown below.

Lattice Analysis					
282.95	143.48	5.24	113.37	288.12	570.57
	435.68	271.68	114.83	5.24	113.37
		463.70	320.28	180.04	61.65
			409.96	269.46	128.65
				303.69	154.29
					164.11

Fig 6.2 Lattice analysis (looking backwards)

By determining whether it is optimal to stay open or close down at each point (whether the expected cash flows from the subsequent period as seen from a point in the lattice is greater than zero), a lattice is developed showing precisely when the option should be exercised. Since it is assumed for these purposes that nothing happens after the end of the lattice period, no strategy is defined in the final column. “NO” indicates it is optimal to keep the system operational, and not to exercise the option. “YES” indicates that closing the system (exercising the option) is optimal.

Strategy - Exercise Option?					
NO	NO	YES	YES	YES	-
	NO	NO	NO	YES	-
		NO	NO	NO	-
			NO	NO	-
				NO	-
					-

Fig 6.3 Lattice analysis strategy

6.2 Inflexible Analysis

An analysis of the system without the put option is performed by multiplying each cash flow by its associated probability of occurrence and summing in each year to give an expected cash flow from that year. The NPV of these expected cash flows is the NPV of the project over six years without the option to close. The figures are again presented in thousands.

Lattice of Cashflows X Probabilities						
	61.13	40.44	2.26	32.00	53.36	69.31
		44.25	46.50	27.38	2.04	36.11
			17.10	29.96	31.48	20.60
				6.28	15.43	22.54
					2.25	7.09
						0.79
Average:	61.13	84.70	61.35	31.62	6.23	54.41
NPV:	174.54					

Fig 6.4 Inflexible system calculation

The value of the option is the difference between the NPVs calculated on the flexible and inflexible bases. In this case, the option is valued at \$108,408. This is a good first order estimate of what the system operator should be willing to pay in order to have the option to close at any time.

As a double-check on this, the backtracking procedure (without the flexibility to close down) is applied to the lattice of cash flows in the inflexible case. The same NPV is obtained as in the previous method of weighing each cash flow by its probability of occurrence and summing.

Backtracking without option					
174.54	32.92	292.04	526.00	663.84	570.57
	425.43	254.49	86.00	53.57	113.37
		463.70	320.28	180.04	61.65
			409.96	269.46	128.65
				303.69	154.29
					164.11

Fig 6.5 Inflexible system calculation using backtracking

6.3 VARG curves from lattice analysis

VARG curves are extracted for both cases – with and without the option to close. This analysis is limited to the first four years of cash flows. Since the outcome of the first year is considered to be deterministic, there are 8 possible paths which can occur over the subsequent periods under consideration. All of these possible paths are enumerated, and their respective probabilities of occurrence determined. The NPV of each path is plotted against cumulative probability to give the VARG curve. The procedure is straightforward in the case where the option to close cannot be executed. In the flexible case which incorporates the option, due care is taken to enumerate only the paths which would be taken given the decision criteria explained in section 6.2 (i.e. That the expected NPV from the following period as seen from the current period exceed zero, which is the cash flow associated with exercising the option and closing down)

One such path is shown below, along with the calculation of the NPV over the first four years. The path in question is that where the price rises in the first two periods after operation begins, and falls during the third period (up, up, down). With reference to the lattice shown the strategy of option execution (Fig 6.3), it can be seen that of the price rises in the first two periods, the optimal strategy is

to exercise the option and shut down the system at the end of the third period. The cash flow for the final period is then zero. The following two lattices show the cash flows for each period, with the extracted values for the path shown in yellow. It should be noted that the “PV cash flows without option” lattice below is equivalent to the early entries of Fig 6.1 expressed in present value terms.

PV Cashflows without option			
61,128	56,046	4,334	85,173
	116,951	85,156	46,319
		119,413	96,654
			115,922
PV Cash flow lattice with closed system			
-	0	0	0
	0	0	0
		0	0
			0

Fig 6.6 Sample path from VARG extraction

The present value of the system of the four periods in this case is $\$61,128 + \$56,046 - \$4,334 + 0 = \$112,840$. This analysis is repeated for all possible paths (there are 8), and the results combined to form the VARG curve. The probability of this path occurring is $p \cdot p \cdot (1-p)$ in the lattice notation (see Section 5).

The VARG curves are shown below. It is important to keep in mind that the curves represent only outcomes that occur over the first four years of cash flows. The flexible case can be seen to remove some of the downside – this corresponds to the case when the price of gas rises so much that the system is closed down early, avoiding losses in the later years. At one point, the flexible case falls behind the fixed design case. This occurs in the case when cash flow from a certain year exceeds the expected cash flow significantly. The decision rule tells the designer to exercise the option and close down the system. The high (but unexpected) cash flow is then foregone in the case of the flexible design. In the event of this high, unlikely cash flow being able to occur as input price falls, the flexible system (which at this point is now closed down) does not benefit.

There is relatively little difference between the two curves, since the decision is generally to keep the system open (to not exercise the option) in the early stages, before the natural gas price has had a chance to rise significantly. If the VARG curves were extended to represent the entire six-year period, a greater divergence would be observed between the two as the option to close is exercised in more and more scenarios as prices rise.

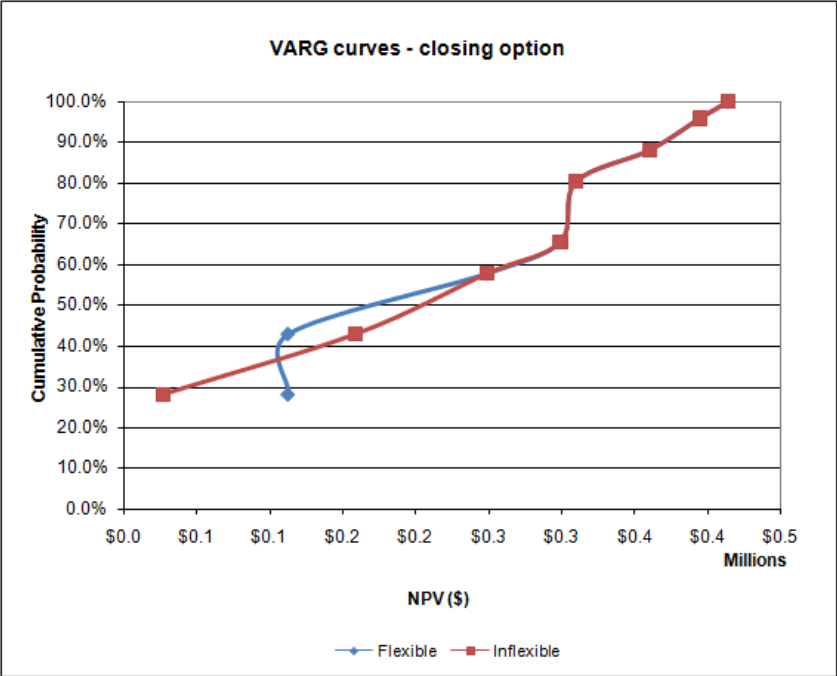


Fig 6.7 VARG curves from lattice analysis

7. Discussion and Conclusions

This portfolio provided a practical and useful means to apply the concepts and tools studied during the semester. The potential value of incorporating flexibility into the design of such a system and the importance of recognising uncertainty were well demonstrated.

As a first consideration, both methods of analysis employed showed an increase in the expected NPV of the system when flexibility was incorporated. As well as this, the VARG curves extracted from the analyses provided a useful means to explore the distribution of the benefits of incorporating flexibility. In both cases, the incorporation of flexibility reduced the system's downside exposure to a greater extent than it increased the upside. A valuable lesson learned is that it is not sufficient to examine the benefits of flexibility in terms of expected outcomes alone. A much more convincing and thorough case can be made to decision makers when the benefit of flexible design (if, indeed, there is a benefit) can be shown across all likely evolutions of uncertainty. This is especially true when one uses a visual representation which is as clear and easily explained as a set of overlaid VARG curves. The table of multiple decision-making criteria produced from the decision analysis was also useful in emphasising the benefit of the flexible case over a broad range of scenarios

The differences between methods of analysis were well demonstrated. As it happened, the method of decision analysis proved better suited to this particular problem. The method proved itself as a more flexible tool in terms of the complexity of path-dependent scenarios that it could consider. A number of simplifying assumptions had to be made in order to conveniently conduct a lattice analysis. Only one uncertainty was considered in the case of the lattice analysis, and a simplified put option on the system was analysed to ensure path independence. In contrast to these shortcomings of the lattice analysis method, its power was readily observed in the high resolution to which the uncertainty of fuel prices could have been analysed, by due to its recombinatorial nature. The computational power of lattice analysis – when it can be appropriately applied – is superior.

A fundamental reality is that uncertain factors which impact on a system can never be accurately forecast. As time passes, there is a distinct benefit associated with being able to make design decisions in a way that takes advantage of increased knowledge as uncertainties are resolved. The expected value of having this design flexibility should comfortably outweigh its cost in order to make it worthwhile. Powerful tools exist which, when used with rational judgement in suitable situations, allow the designer to calculate the value of real options in design.

8. Appendices

8.1 Simple explanation of boiler efficiency theory used

The *part load ratio* (PLR) of a boiler is defined as

$$PLR = \frac{Q_o}{Q_{o,full}}$$

Where Q_o is the heat output from the boiler at a given time and $Q_{o,full}$ is the maximum possible heat output – the rated capacity – of the boiler

The energy input (obtained from burning natural gas) required to achieve this level output is given by a polynomial equation of the form

$$\frac{Q_i}{Q_{i,full}} = A + B(PLR) + C(PLR)^2 + \dots$$

Where A, B, C... are constants which characterise the particular boiler. In practice this equation is truncated after the squared term.

The rated capacity of the boiler and the efficiency of the boiler when operating at full capacity are specified. This efficiency, η , is given by the relation

$$\eta = \frac{Q_{o,full}}{Q_{i,full}}$$

This is used to determine the heat input required when the boiler operates at capacity. The above polynomial can then be used to calculate the heat input required to meet various heating loads, and hence to specify the efficiency of the boiler, which varies with the heating load.