

Introduction to Real Options

We introduce real options and discuss some of the issues and solution methods that arise when tackling these problems. Our main example is the Simplico gold mine example from Luenberger. This contains many of the features typically found in real options applications – a non-financial setting, some financial uncertainty and an element of control. An important feature that is not contained in the Simplico example is non-financial uncertainty but other examples will include this. Real options analysis is an important aspect of *capital budgeting*, one of the central problems of corporate finance, where the goal is to figure out the fair risk-adjusted value of investment opportunities and to then decide which opportunities – if any – should be pursued.

1 Introduction to Real Options

The term "real options" is often used to describe investment situations involving non-financial, i.e. *real*, assets together with some degree of *optionality*. For example, an industrialist who owns a factory with excess capacity has an option to increase production that she may exercise at any time. This option might be of particular value when demand for the factory's output increases. The owner of an oilfield has an option to drill for oil that he may exercise at any time. In fact since he can drill for oil in each time period he actually holds an entire series of options. On the other hand if he only holds a lease on the oilfield that expires on a specified date, then he holds only a finite number of drilling options. As a final example, consider a company that is considering investing in a new technology. If we ignore the optionality in this investment, then it may have a negative NPV so that it does not appear to be worth pursuing. However, it may be the case that investing in this new technology affords the company the option to develop more advanced and profitable technology at a later date. As a result, the investment might ultimately be a positive NPV value project that is indeed worth pursuing.

In this section we describe some of the main ideas and issues that arise in the context of real options problems. It is perhaps unfortunate that the term "real options" has become such a buzzword in recent years, as this has helped fuel the exaggerated claims that are often made by 'practitioners' on its behalf. For example, proponents often claim[1] that the NPV investment criterion is flawed as it does not take optionality into account. This claim makes little sense, however, as there is no reason why a good NPV analysis would not account for optionality. The above criticisms notwithstanding, it is worth stating that real options problems arise in many important contexts, and that economics and financial engineering provide useful guidance towards solving these problems. The principal characteristics shared by real options problems are:

1. They involve non-financial assets, e.g. factory capacity, oil leases, commodities, technology from R&D etc. It is often the case, however, that financial variables such as foreign exchange rates, commodity prices, market indices etc. are also (sometimes implicitly) present in the problem formulation.

2. The natural framework for these problems is one of incomplete markets as the stochastic processes driving the non-financial variables will not be *spanned*[2] by the financial assets. For example, it is not possible to construct a self-financing trading strategy in the financial markets that replicates a payoff whose value depends on whether or not there is oil in a particular oilfield, or whether or not a particular manufacturing product will be popular with consumers. Because of the market incompleteness, we do not have a unique

[1]Some knowledgeable practitioners also make this claim, choosing to interpret the phrase 'NPV analysis' as 'naive NPV analysis' that does not take optionality into account. I can see no justification for this!

[2]That is, a security that depends at least in part on such a stochastic process will not be attainable using a self-financing trading strategy in the financial market.

EMM that we can use to evaluate the investment opportunities. We therefore use (financial) economics theory to guide us in choosing a good EMM (or indeed a good set of EMM's) that we should work with. It is worth remarking again that choosing an EMM is equivalent to choosing a (stochastic) discount factor. We can therefore see that the problem of choosing an EMM may be formulated in the traditional economics and corporate finance context as a problem of finding the correct way to discount cash-flows.

3. There are usually options available to the decision-maker. More generally, real options problems are usually *control* problems where the decision-maker can (partially) control some of the quantities under consideration. Moreover, it is often inconvenient or unnecessary to explicitly identify all of the 'options' that are available. For example, in the Simplico gold mine example below, the owner of the lease has an option each year to mine for gold but in valuing the lease we don't explicitly consider these options and evaluate them separately.

Our first example of a real options problem concerns valuing a lease on a gold mine and an option to increase the rate at which gold can be extracted.

Example 1 (**Luenberger's** Simplico Gold Mine)

Gold Price Lattice

											2476.7
										2063.9	1857.5
									1719.9	1547.9	1393.1
								1433.3	1289.9	1161.0	1044.9
							1194.4	1075.0	967.5	870.7	783.6
						995.3	895.8	806.2	725.6	653.0	587.7
					829.4	746.5	671.8	604.7	544.2	489.8	440.8
				691.2	622.1	559.9	503.9	453.5	408.1	367.3	330.6
			576.0	518.4	466.6	419.9	377.9	340.1	306.1	275.5	247.9
		480.0	432.0	388.8	349.9	314.9	283.4	255.1	229.6	206.6	186.0
	400.0	360.0	324.0	291.6	262.4	236.2	212.6	191.3	172.2	155.0	139.5
Date	0	1	2	3	4	5	6	7	8	9	10

Gold can be extracted from the simplico gold mine at a rate of up to $10,000$ ounces per year at a cost of $C = \$200$ per ounce. The current market price of gold is $\$400$ and it fluctuates randomly in such a way that it increases each year by a factor of 1.2 with probability $.75$ or it decreases by a factor of $.9$ with probability $.25$, i.e. gold price fluctuations are described by a binomial model with each period corresponding to 1 year. Interest rates are flat at $r = 10\%$ per year. We want to compute the price of a lease on the gold mine that expires after 10 years. It is assumed that any gold that is extracted in a given year is sold at the end of the year at the price that prevailed at the beginning of the year. The gold price lattice is shown above.

Since gold is a traded commodity with liquid markets we can obtain a unique risk-neutral price for any derivative security dependent upon its price process. The risk-neutral probabilities are found to be $q = 2/3$ and $1 - q = 1/3$. The price of the lease is then computed by working backwards in the lattice below. Because the lease expires worthless the node values at $t = 10$ are all zero. The value 16.9 on the uppermost node at $t = 9$, for example, is obtained by discounting the profits earned at $t = 10$ back to the beginning of the year. We therefore obtain $16.94 = 10,000(2063.9 - 200)/1.1$. The value at a node in any earlier year is obtained by discounting the value of the lease *and* the profit obtained at the end of the year back to the beginning of the year and adding the two quantities together. For example, in year 6 the central node has a value of 12 million and this is obtained as

$$12,000,000 \;=\; \frac{10,000(503.9 - 200)}{1.1} \;+\; \frac{q \times 11,500,000 + (1-q) \times 7,400,000}{1.1}.$$

Note that in any year when the price of gold is less than $\$200$, it is optimal to extract no gold and so no profits are recorded for that year. This simply reflects the fact that every year the owner of the lease has an option to extract gold or not, and will do so only when it is profitable to do so. Backwards evaluation therefore takes the form

$$V_t(s) = \frac{10\text{k} \times \max\{0, s - C\} + (qV_{t+1}(us) + (1-q)V_{t+1}(ds))}{1+r}$$

where $V_t(s)$ is the value of the lease at time t when the gold price is s. We find the value of the lease at $t = 0$ is 24.1 million.

Lease Value (in millions)

											0.0
										16.9	0.0
									27.8	12.3	0.0
								34.1	20.0	8.7	0.0
							37.1	24.3	14.1	6.1	0.0
						37.7	26.2	17.0	9.7	4.1	0.0
					36.5	26.4	18.1	11.5	6.4	2.6	0.0
				34.2	25.2	17.9	12.0	7.4	3.9	1.5	0.0
			31.2	23.3	16.7	11.5	7.4	4.3	2.1	0.7	0.0
		27.8	20.7	15.0	10.4	6.7	4.0	2.0	0.7	0.1	0.0
	24.1	17.9	12.9	8.8	5.6	3.2	1.4	0.4	0.0	0.0	0.0
Date	0	1	2	3	4	5	6	7	8	9	10

Suppose now that it is possible to enhance the extraction rate to $12,500$ ounces per year by purchasing new equipment that costs $\$4$ million. Once the new equipment is in place then it remains in place for all future years. Moreover the extraction cost would also increase to $\$240$ per year with the enhancement in place and at the end of the lease the new equipment becomes the property of the original owner of the mine. The owner of the lease therefore has an *option* to install the new equipment at any time and we wish to determine the value of this option. To do this, we first compute the value of the lease assuming that the new equipment is in place at $t = 0$. This is done in exactly the same manner as before and the values at each node and period are given in the following lattice:

Lease Value Assuming Enhancement in Place (in millions)

											0.0
										20.7	0.0
									33.9	14.9	0.0
								41.4	24.1	10.5	0.0
							44.8	29.2	16.8	7.2	0.0
						45.2	31.2	20.0	11.3	4.7	0.0
					43.5	31.0	21.0	13.2	7.2	2.8	0.0
				40.4	29.3	20.4	13.4	8.0	4.1	1.4	0.0
			36.4	26.6	18.7	12.5	7.7	4.1	1.8	0.4	0.0
		31.8	23.3	16.3	10.8	6.5	3.4	1.3	0.2	0.0	0.0
	27.0	19.5	13.5	8.6	4.9	2.3	0.8	0.1	0.0	0.0	0.0
Date	0	1	2	3	4	5	6	7	8	9	10

We see that the value of the lease at $t = 0$ assuming that the new equipment is in place (the $\$4$ million cost of the new equipment has not been subtracted) is $\$27$ million. We let $V_t^{\text{eq}}(s)$ denote the value of the lease at time t when the gold price is s and the new equipment is in place.

We now value the *option* to install the new equipment as follows. We construct yet another lattice (shown below) that, starting at $t = 10$, assumes the new equipment is not in place, i.e. we begin by assuming the *original* equipment and parameters apply. We work backwards in the lattice, computing the value of the lease at each node as before but now with one added complication: after computing the value of the lease at a node we compare this value, A say, to the value, B say, at the corresponding node in the lattice where the enhancement was assumed to be in place. If $B - \$4\text{million} \geq A$ then it is optimal to install the equipment at this node, if it has not already been installed. We also place $\max(B - \$4m,\ A)$ at the node in our new lattice. More specifically, let us use $U_t(s)$ to denote the price of the lease at time t when the gold price is s and with the option to enhance in place. Then we can solve for $U_t(s)$ as follows:

$$U_t(s) = \max\left\{V_t^{\text{eq}}(s) - 4\text{m},\ \frac{10\text{k} \times \max\{0, s - C\} + (qU_{t+1}(us) + (1-q)U_{t+1}(ds))}{1+r}\right\}. \tag{1}$$

We continue working backwards using (1), determining at each node whether or not the new equipment should be installed if it hasn't been installed already. The new lattice of $U_t(s)$ values is displayed below.

Lease Value with Option for Enhancement (in millions)

											0.0
										16.9	0.0
									29.9*	12.3	0.0
								37.4*	20.1*	8.7	0.0
							40.8*	25.2*	14.1	6.1	0.0
						41.2*	27.2*	17.0	9.7	4.1	0.0
					39.5*	27.0*	18.1	11.5	6.4	2.6	0.0
				36.4*	25.6	17.9	12.0	7.4	3.9	1.5	0.0
			32.6	23.5	16.7	11.5	7.4	4.3	2.1	0.7	0.0
		28.6	20.9	15.0	10.4	6.7	4.0	2.0	0.7	0.1	0.0
	24.6	18.0	12.9	8.8	5.6	3.2	1.4	0.4	0.0	0.0	0.0
Date	0	1	2	3	4	5	6	7	8	9	10

We see that the value of the lease with the option is $\$24.6$ million, slightly greater than the value of the lease without the option. Entries in the lattice marked with a '*' denote nodes where it is optimal to install the new equipment. (Note that the values at these nodes are the same as the values at the corresponding nodes in the preceding lattice less $\$4$ million.) ∎

Example 1 is an interesting real options problem as there is no non-financial noise in the problem, implying in particular that the market is complete. As mentioned earlier, most real options problems have non-financial noise and therefore require an *incomplete markets* treatment. Example 1 could easily be generalized in this direction by introducing uncertainty regarding the amount of gold that can be extracted or by introducing stochastic storage or extraction costs for the gold. Note that for these generalizations the financial component of the problem, i.e. the binomial lattice for gold prices, would still be complete and so we would still have unique risk-neutral probabilities for the up- and down-moves in the price of gold. We would, however, need to use other economic considerations to help us determine[3] the risk-neutral dynamics of storage and extraction costs etc. When there is non-financial noise and the problem is therefore incomplete, it is often a good modelling technique to simply assume the financial market in your model is complete so that any random variable, i.e. contingent claim, that depends *only* on financial noise can indeed be replicated by a self-financing trading strategy in the financial assets.

[3]Note that a storage cost process or extraction rate process is not a financial asset and so martingale pricing theory does not imply that these processes should be martingales under an EMM.

Exercise 1 *Compute the value of the enhancement option of Example 1 when the enhancement costs* $\$5$ *million but raises the mine capability by* 40% *to* $14,000$ *ounces at an operating cost of* $\$240$ *per ounce. Moreover, due to technological considerations, you should assume that the enhancement (should it be required) will not be available until the beginning of the* 5^{th} *year.*

---

2 Pricing Methods for Real Options Problems

In Example 1 the market was complete and so there was only one valuation technique available, i.e. martingale pricing using the *unique* equivalent martingale measure. When markets are incomplete there are many pricing methods available. They include:

1. **Zero-Level Pricing**: Zero-level pricing is based on utility maximization for portfolio optimization problems. It is that security price that leaves the decision-maker indifferent between purchasing and not purchasing an *infinitesimal* amount of the security.
2. **Utility Indifference Pricing**: If the decision-maker is risk-averse and has a utility function then the problem may be recast as a portfolio optimization problem where the price of the real option is taken to be that price that leaves the agent indifferent between purchasing and not purchasing the *entire* project (while simultaneously allowing dynamic trading in the financial assets). This method might be suitable for a risk averse decision maker considering a project that could only be undertaken at the zero-level or in its entirety.
3. **Super-Replication and Quadratic Hedging**: These pricing methods are based on *hedging* cash-flows.
4. **'Good-Deal Bounds'**: This is a heuristic method based on the observation that *expected* returns on any security should be neither too 'high' nor too 'low'. The method simultaneously considers many plausible EMMs and constructs an interval of plausible prices for the real option.
5. **Projection Methods**: These methods work by projecting the real option's cash flows onto the space spanned by the financial assets. An EMM that prices this projection correctly is then used with the remaining part of the cash-flow priced using the true probability distribution and discounting with the cash account.

In general, these methods may be interpreted as selecting a particular EMM for pricing the cash-flow in question, and so they are consistent with martingale pricing theory. The 'good-deal bounds' method might also be interpreted as a *robust* pricing method as it (effectively) considers many plausible EMMs and constructs an *interval* of plausible prices.

Zero-Level Pricing with Private Uncertainty

We say that a source of uncertainty is *private* if it is independent of any uncertainty driving the financial markets and is specific to the problem at hand. For example, the success of an R&D project, the quantity of oil in an oilfield, the reliability of a vital piece of manufacturing equipment or the successful launch of a new product are all possible sources of private uncertainty. More generally, while the incidence of natural disasters or political upheavals would not constitute sources of private uncertainty, it would be very difficult[4], if not impossible, to adequately hedge against these events using the financial markets. As a result, they could often be treated as if they did constitute private sources of uncertainty.

Economic considerations suggest that if we want to use zero-level pricing to compute real option prices when there is only private uncertainty involved, then we should use the true probability distribution to do so and

[4]As markets become ever more sophisticated, however, it should become increasingly possible to hedge, at least partly, against some of these events. For example, the introduction of catastrophe bonds and the expanding use of financial technology in the insurance industry suggest that we will be able to hedge against ever more sources of uncertainty in the future.

discount by the risk-free interest rate. A CAPM-based argument can provide some intuition for this observation. In particular, the CAPM states

$$\mathrm{E}[r_o] \;=\; r_f + \beta_o \left(\mathrm{E}[r_m] - r_f\right)$$

where $\beta_o := \mathrm{Cov}(r_m, r_o)/\mathrm{Var}(r_m)$ and r_m is the return on the market portfolio. Therefore, if r_o is the return on an investment that is only exposed to private uncertainty so that $\mathrm{Cov}(r_m, r_o) = 0$, the CAPM implies $\mathrm{E}[r_o] = r_f$. This implies[5] the value, P_0, of the investment in a CAPM world is given by

$$P_0 \;=\; \frac{\mathrm{E}[P_1]}{1 + r_f}$$

where P_1 is the terminal payoff of the investment. This and other similar arguments have been used to motivate the practice of using risk-neutral probabilities to price the 'financial uncertainty' of an investment and the true probabilities to price the 'non-financial uncertainty' of the investment. Note that this practice is consistent with martingale pricing as the use of risk-neutral probabilities ensures that deflated (financial) security prices are martingales. In particular, this methodology ensures that there are no arbitrage opportunities available if trading in the financial markets is also permitted as part of the problem formulation.

This practice (of using the true probability distribution to price private uncertainty) is much easier to justify if we can *disentangle* the financial uncertainty from the non-financial uncertainty. For example, consider a one period model where at $t = 1$ you will have X barrels of oil that you can then sell in the spot oil markets for $\$P$ per barrel. You therefore will earn XP at $t = 1$. If X is a random variable that is only revealed to you at $t = 1$, then it is impossible to fully hedge your oil exposure by trading at $t = 0$ as you do not know then how much oil you will have at $t = 1$. On the other hand if X is revealed to you at $t = 0$ then you can fully hedge at $t = 0$ your resulting oil exposure using the 1-period oil futures. Only non-financial uncertainty (which is revealed at $t = 0$) then remains and so the situation is similar to that in our CAPM argument above. The reasons for using the true probabilities to price the non-financial uncertainty (X in our example) is therefore more persuasive here. In the former case where X is not revealed until $t = 1$, the financial uncertainty cannot be perfectly hedged and the economic argument for using the true probabilities to price the non-financial uncertainty is not as powerful. Still it can be justified by considering, for example, a first order Taylor expansion of the payoff, XP, about the means, $\bar{X}$ and $\bar{P}$, to obtain

$$\begin{aligned} XP &\approx \bar{X}\bar{P} + \bar{P}(X - \bar{X}) + \bar{X}(P - \bar{P}) \\ &= \bar{P}X + \bar{X}P - \bar{X}\bar{P}. \end{aligned} \tag{2}$$

Note that the payoff approximation (2) does allow for the financial and non-financial uncertainties to be disentangled. The above CAPM argument can now be used to price the non-financial uncertainty, $\bar{P}X$, using the true probability distribution.

2.1 Real Options in Practice

In practice, real options problems are often too complex to be solved exactly, either analytically or numerically. There are many reasons for this:

1. High dimensionality due to the presence of many state variables, control variables and / or sources of uncertainty.
2. Complexity of real world constraints. For example, in Example 1 we assumed that the mine operator could choose not to extract gold when prices were unfavorable. However, political considerations might not allow such action to be taken as it would presumably result in the unemployment of most people associated with operating the mine.
3. Data uncertainty. Every model is limited by the quality of the data and in particular, the quality of parameter estimates. Real options problems are no different in this regard.
4. Game theoretic considerations. Sometimes it might be necessary take the actions of competitors into account but this only complicates the analysis further.

[5] This is clear if we write $P_1 = (1 + r_o)P_0$ and take expectations.

Because of these complications, it often makes more sense to seek good approximate solutions to real options problems. There is a tradeoff between model complexity and tractability and finding the right balance between the two is more of an art than a science. While there has not been a lot of (academic) work done in this area, techniques such as simulation and approximate dynamic programming methods should prove useful. Moreover, if the solution technique is utility independent, e.g zero-level pricing, then economic considerations should be used to guide the choice of the EMM or equivalently, the stochastic discount factor. More generally, by considering several 'plausible' EMMs it should be possible to construct an *interval* of plausible valuations. Such an interval might be more than adequate in many contexts.

2.2 The Incorrectness of Using Deterministic Discount Factors

In practice, however, corporations often (incorrectly) use a single deterministic discount factor to value cash-flows. While the corporate finance literature has long recognized that this is not appropriate, whether or not the resulting valuations are significantly skewed will depend on the problem in question.

In order to (re)emphasize the necessity of using a stochastic discount factor, consider the Black-Scholes framework where the stock-price follows a geometric Brownian motion and there is a constant interest rate, r. In this model the stock price at time T is given by

$$S_T = S_0 e^{(\mu-\sigma^2/2)T+\sigma B_T}$$

where B_T is the time T value of a standard Brownian motion. Note that the only deterministic discount factor that correctly prices the stock is $d = \exp(-\mu T)$. However using this discount factor to price options on the stock is *incorrect* and will produce arbitrage opportunities! In fact the only way to price all securities depending on the underlying Brownian motion correctly is to use a *stochastic* discount factor[6] or equivalently, an equivalent martingale measure. You can easily check that the same comments apply to pricing in the discrete-time models we have considered in this course. That said, if a particular payoff is very similar to S_T then using the deterministic discount factor, $d = \exp(-\mu T)$, to price the payoff should lead to a reasonably accurate result.

---

3 Some Further Examples

We finish with a couple of further examples on pricing real options. The first example applies zero-level pricing to compute the value of a foreign venture whereas the second example requires a dynamic programming formulation as an optimal control strategy is required to figure out the value of the investment.

Example 2 (Valuing A Foreign Venture)
A particular investment gives you the rights to the monthly profits of a foreign venture for a fixed period of time. The first payment will be made one month from now and the final payment will be in 5 months time after which the investment will be worthless. The monthly payments are denominated in *Euro*, and are IID random variables with expectation μ. They are also independent of returns in both the domestic and foreign financial markets. You would like to determine the value of the investment.

Let us first assume that the domestic, i.e. US, interest rate is 5% per annum, compounded monthly. This implies a gross rate of 1.0042 per month. Similarly, the annual interest rate in the Euro zone, i.e. the foreign interest rate, is assumed to be 10%, again compounded monthly. This implies a per month gross interest rate of 1.0083. We can construct a binomial lattice for the $/Euro exchange rate process if we view the foreign currency, i.e. the Euro, as an asset that pays dividends, i.e. interest, each period. Martingale pricing in a binomial model for the exchange rate with up- and down-factors, u and d respectively, implies that

$$X_i = \mathrm{E}_i^Q \left[\frac{X_{i+1} + r_f X_i}{1 + r_d} \right] \qquad (3)$$

[6] In fact the early research into option pricing was focussed on identifying the correct discount factor.

where X_i is the \$/Euro exchange rate at time i, r_d and r_f are the domestic and foreign per-period interest rates, respectively, and the risk-neutral probability, q, of an up-move satisfies

$$q = \frac{1 + r_d - d - r_f}{u - d}. \tag{4}$$

The binomial lattice is given below with $X_0 = 1.20$, $u = 1.05$ and $d = 1/u$.

Dollar/Euro Exchange Rate					
					1.53
				1.46	1.39
			1.39	1.32	1.26
		1.32	1.26	1.20	1.14
	1.26	1.20	1.14	1.09	1.04
1.20	1.14	1.09	1.04	0.99	0.94
t=0	t=1	t=2	t=3	t=4	t=5

Valuing the investment using zero-level pricing (and therefore using the the true probability measure, P, for the non-financial uncertainty) is now straightforward. At each time-t node in the lattice we assume there is a cash-flow of μX_t. These cash-flows are valued as usual by backwards evaluation using the risk-neutral probabilities computed in (4). Can you see an easy way (that does not require backwards evaluation) to value the cash-flows? ∎

Exercise 2 *If the monthly cash-flows and the foreign financial market were* dependent, *how would you go about valuing the security? What assumptions would you need to make?*

Example 19.5 of Luenberger (2^{nd} ed.) provides another more complex example of zero-level pricing in the context of valuing the so-called *rapido* oil well. In that example the market uncertainty is the price of oil and risk-neutral probabilities are used to "value" that uncertainty. In contrast, true probabilities are used to "value" the private uncertainty which relates to the initial flow of oil and how quickly that flow reduces over time.

Example 3 (Operating a Gas-Fired Power Plant)
A power company is considering renting a two-regime gas-fired power plant over the (discrete) time $[0, T]$. In order to figure out whether or not this is worthwhile, the company must figure out the optimal operating policy for the plant. At any time t the company can operate the plant in "shut-down" mode, low-capacity mode or high-capacity mode. The relevant details for each of these three operating states are:

- Shut-down mode: plant rent K
- Low capacity mode: output $\underline{Q}$, gas consumption $\underline{H}$, plant rent $\underline{K}$
- High capacity mode: output $\overline{Q}$, gas consumption $\overline{H}$, plant rent $\overline{K}$

If the plant is currently in operation then shutting the plant down will incur a ramp-down cost of C_d. Similarly, if the plan is currently shut down, then moving it into a low / high production mode will incur a ramp-up cost of C_u. We use $s_t \in \{0 \equiv \text{plant shut-down},\ 1 \equiv \text{plant operating}\}$ to denote the time t state of the plant. In each period there are two possible actions $a \in \{0 \equiv \text{operate in shut-down mode},\ 1 \equiv \text{operate in production mode}\}$.

We let $V_t(s, P, G)$ denote the optimal profit of the plant over the horizon $[t, T]$ when the current state is s, the current price of electricity is P, and the current price of gas is G. We also have the following additional notation:

- $c(s, a)$: cost of taking action a when the plant is in state s
- $u(s, a)$: new state of the plant when action a is taken in state s

Specific values are:

$$
\begin{array}{rcl}
c(0,0) & = & K \\
c(0,1) & = & C_u + K \\
c(1,0) & = & C_d + K \\
c(1,1) & = & \min\left\{\overline{K} + G_t\overline{H} - P_t\overline{Q},\ \underline{K} + G_t\underline{H} - P_t\underline{Q}\right\}
\end{array}
\qquad
\begin{array}{rcl}
u(0,0) & = & 0 \\
u(0,1) & = & 1 \\
u(1,0) & = & 0 \\
u(1,1) & = & 1
\end{array}
$$

Note that when the plant is in production mode one has the choice to run it at the low or high capacity mode. We can now write the standard dynamic programming iteration:

$$V_t(s_t, P_t, G_t) = \max_{a \in \{0,1\}} \left\{ -c(s_t, a) + e^{-r\Delta t} \mathrm{E}_t^Q \left[V_{t+1}(u(s_t, a), P_{t+1}, G_{t+1}) \right] \right\} \tag{5}$$

where $\mathrm{E}_t^Q[\cdot]$ denotes the usual risk-neutral expectation conditional on all known-information at time t, r is the continuously compounded risk-free rate and Δt is the length of a period. There are two sources of uncertainty here: the price dynamics of electricity and gas. It is straightforward to construct lattices for these prices and solve (5) numerically beginning at time T and working backwards. ∎

Exercise 3 *How should we choose the appropriate EMM, Q, to use in (5)?*