IEOR E4706: Financial Engineering: Discrete-Time Asset Pricing
© 2010 by Martin Haugh

Martingale Pricing Applied to Dynamic Portfolio Optimization and Real Options

We consider some further applications of martingale pricing to problems in financial engineering. In particular, we will show how dynamic portfolio optimization problems in complete markets may be solved using martingale pricing methods. We will see, as a result, how the problems of security pricing and portfolio optimization are very closely related. We then introduce real options and discuss some of the issues and solution methods that arise when tackling these problems.

Portfolio Optimization: The Martingale Method

The traditional technique for solving dynamic portfolio optimization problems is dynamic programming (DP), or as it is known in continuous-time applications, control theory. While DP is a very important technique, it suffers from the so-called *curse of dimensionality*. That is, the computational requirements of a DP (that needs to be solved numerically) grow exponentially in the dimensionality of the problem. Any new techniques for solving dynamic control problems are therefore most welcome. One such technique that applies to portfolio optimization problems in complete markets, is the *martingale method*.

The idea behind this method is very simple: since markets are complete, every contingent claim is attainable and therefore the initial price of any random variable representing a contingent claim can be computed. Now consider a portfolio optimization problem where, for example, we seek to maximize expected utility of terminal wealth, $\mathrm{E}[U(W_T)]$. Instead of trying to solve for the optimal dynamic trading strategy, we simply choose the optimal (contingent claim) W_T^* subject to a budget constraint stating that the cost of W_T^* may not exceed the initial wealth, W_0. Once W_T^* has been chosen, we then find the trading strategy, θ_t^*, that replicates W_T^*. The problem of finding W_T^* is therefore a *static* optimization problem that can be decomposed into two stages: find W_T^* and then find θ_t^*. In comparison, the problem of directly finding the optimal trading strategy is a dynamic problem and in general more difficult to solve. Consider the following example.

Example 1 (Portfolio Optimization: Power Utility)

Consider the complete market model of Example 9 in the *Martingale Pricing Theory* lecture notes. The evolution of information is represented in the tree below and the prices of each of the three securities are given at each node.

Let p_i denote the true probability that the terminal state will be ω_i and let $\pi_0^{\{2\}}(i) := \pi_0^{\{2\}}(\omega_i)$ represent[1] the i^{th} state price as described earlier in the *Martingale Pricing Theory* lecture notes. Suppose now that you wish to solve the following portfolio optimization problem where we wish to maximize the expected utility of terminal wealth:

$$\max_{W_T} \mathrm{E}\left[U(W_T)\right] \qquad \text{subject to} \qquad \sum_{i=1}^{9} \pi_0^{\{2\}}(i) W_{T,i} = W_0.$$

Instead of using dynamic programming, we will solve the problem using the martingale method.

Note that because there are nine possible states, the problem amounts to choosing $W_T := (W_{T,1}, \ldots, W_{T,9})$ where $W_{T,i} := W_T(\omega_i)$. If we have a *power* utility function, i.e. $U(W_T) = W_T^{1-\gamma}/(1-\gamma)$ for $\gamma > 0$, then the

[1] That is, $\pi_0^{\{2\}}(i)$ is the date 0 price of a security that pays \$1 at date 2 if and only if the true state is ω_i.

problem may be expressed as

$$\max_{W_{T,i}} \sum_{i=1}^{9} p_i \frac{W_{T,i}^{1-\gamma}}{1-\gamma} \qquad \text{subject to} \qquad \sum_{i=1}^{9} \pi_0^{\{2\}}(i) W_{T,i} = W_0. \tag{1}$$

This is a static optimization problem and may be solved by using a single Lagrange multiplier, λ, for the budget constraint.

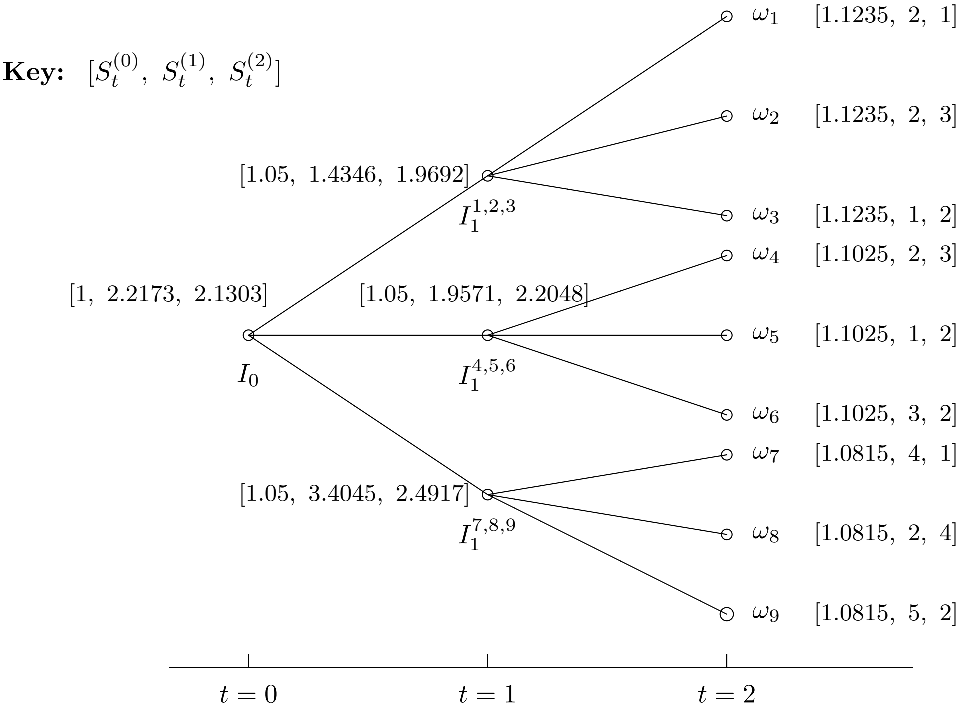

The first order conditions are

$$p_i W_{T,i}^{-\gamma} \; - \; \lambda \pi_0^{\{2\}}(i) \; = \; 0 \quad \text{for} \;\; i = 1, \ldots, 9. \tag{2}$$

In particular, this implies

$$W_{T,i} \;=\; \left(\frac{\lambda \pi_0^{\{2\}}(i)}{p_i} \right)^{-1/\gamma} \tag{3}$$

and we can now use the budget constraint to compute

$$\lambda^{* \, -1/\gamma} \;=\; \frac{W_0}{\sum_{i=1}^{9} \pi_0^{\{2\}}(i)^{(\gamma-1)/\gamma} \;\; p_i^{1/\gamma}}. \tag{4}$$

We can now substitute (4) into (3) to obtain the optimal wealth, W_T^*. We can also obtain that the optimal objective function is given by

$$\begin{aligned} \mathrm{E}\left[U(W_T^*)\right] \;&=\; \mathrm{E}\left[\frac{W_T^{1-\gamma}}{1-\gamma}\right] \\ &=\; \frac{W_0^{1-\gamma}}{1-\gamma} \; \mathrm{E}\left[\Pi_0^{\{2\}\,(\gamma-1)/\gamma}\right]^{\gamma} \end{aligned} \tag{5}$$

where $\Pi_0^{\{2\}}$ is the state-price density (SPD) or *pricing kernel*. The SPD is simply the vector of state prices divided by the *true* probability of each state, i.e., $\Pi_0^{\{2\}}(\omega_k) = \pi_0^{\{2\}}(k)/p_k$.

The final stage of the solution is to find the trading strategy that replicates W_T^*. This is straightforward and is done by working backwards in time, starting from date $t = 2$. ∎

Remarks

1. The martingale technique can be applied to a broad range of complete-market problems in both discrete- and continuous-time models. For example, it is easy to handle the case where utility is obtained in each period t by consuming some of the wealth, W_t.

2. Note that the method will not work if the market is incomplete. This is because we do not have a unique set of state prices and so we do not know how to price all the possible contingent claims, W_T. In particular, it is not the case that every possible W_T is attainable. It is also worth remarking at this point that trading constraints e.g. no-borrowing or no short-selling, can also render a market incomplete.

3. In recent years, however, there has been a considerable amount of research directed at extending the martingale method to incomplete market problems. This effort has been quite successful and a number of related methods have been developed. One such method, for example, uses duality theory to identify an optimal set of state prices that should be used.

4. At the beginning of the course we mentioned that the problems of security pricing and portfolio optimization are intimately related. We now see why this is the case as both problems rely on the availability of state prices, or equivalently, martingale probabilities.

The Dynamic Programming (DP) Approach to our Portfolio Optimization Problem

We could also attempt to solve portfolio optimization problems using dynamic programming. Assume again that we wish to maximize the expected utility of terminal wealth, $\mathrm{E}[u(W_T)]$. We write $V_t(W, I)$ for the value function at time t when the current wealth is W and we are at node I in the tree. For example, in the notation of Example 1 we therefore want to find $V := V_0(W_0, I_0)$. If there are N periods, R_j is the vector of gross returns on the securities between dates j and $j+1$, and θ_j is the corresponding vector of portfolio *weights*, then our DP algorithm is:

> **set** $V_N(W, I) := u(W)$ for all W and terminal nodes, I
> **for** $j = N-1$ down to 0
>
> **Solve** for all W and date j nodes, I:
>
> $$V_j(W, I) := \max_{\theta_j} \ \mathrm{E}_j \left[V_{j+1}(W \theta_j^T R_j, \mathsf{succ}(I)) \right]$$
>
> **set** $V := V_0(W_0, I_0)$

where succ(I) in the above algorithm refers to the time $j+1$ node that succeeds the time j node, I.

Remark 1 *The difficulty with using DP in general is that it is typically the case that the maximization problem in the above algorithm cannot be solved explicitly. The DP problem then needs to be solved numerically, thereby limiting the size of problems that can be practically solved. Note also that when using the DP approach we need to find the date t value function, $V_t(\cdot, \cdot)$, for all possible values of W. The martingale method did not need to do this, requiring only a single optimization problem for a single wealth value, W_0.*

It so happens that for the case of power utility (as in Example 1) it may be shown that the maximization problem in the DP algorithm has the same[2] solution, θ_j^*, regardless of the value of W. That is, the optimal

[2] This only holds when the problem is formulated so that θ represents the vector of portfolio weights, that is, the fractions of wealth that are invested in the securities. (With this interpretation, it is clear that the portfolio weights must sum to one.)

trading strategy is wealth independent. The DP approach in this case therefore turns out to be more straightforward than usual, but it is still more difficult to implement that the martingale approach.

Exercise 1 *Show that the optimal θ_j^* in Example 1 is wealth independent.*

The State-Price-Density Process

We have seen how the concept of state prices are important for financial engineering problems. As we mentioned in Example 1, we can equivalently work with the *state price density*, Π, where $\Pi(\omega)$ is defined to be $\pi(\omega)/p(\omega)$, i.e., the state price divided by the probability of the state.

As time elapses in dynamic models, it is clear that the probabilities of the various states and their state prices should adapt to the new information that arises. Just as security prices vary stochastically, so too do the state prices and state probabilities. In particular, just as security prices are stochastic processes, so too is the state price density a stochastic process. We refer to it simply as the SPD process.

Example 2 (The SPD Process)

Recall Example 9 in the *Martingale Pricing Theory* lecture notes where we computed the date $t = 0$ state prices, $\pi_0^{\{2\}}(\omega_i)$ for $i = 1, \ldots, 9$. We also computed, for example, the state prices, $\pi_1^{\{2\}}(\omega_i)$ and found that

$$\pi_1^{\{2\}}(\omega_1) = .2 \text{ , } \pi_1^{\{2\}}(\omega_2) = .3 \text{ and } \pi_1^{\{2\}}(\omega_3) = 0.4346 \text{ at node } I_1^{1,2,3}.$$

What we didn't say (but was implicitly understood) was that $\pi_1^{\{2\}}(\omega_i) = 0$ at node $I_1^{1,2,3}$ for $4 \leq i \leq 9$. In other words, we have seen how the values of the nine state prices depend on the date and the particular node at which we find ourselves. That is, the state prices, and therefore the SPD, constitute a stochastic process. ∎

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. For a final example, consider a company that is considering investing in a new technology. Taken in isolation, this investment 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[3] 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. A more accurate term than "real options" might be "investment under uncertainty" but this is not quite glamorous enough so its usage[4] appears to be confined to the academic literature in finance and economics.

The above criticisms notwithstanding, it is worth stating that real options[5] problems arise in many important

[3]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!

[4]Indeed, the standard academic reference for real options is *Investment Under Uncertainty* by Dixit and Pindyck, a PhD level economics text.

[5]We will persevere with the term "real options".

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, interest rates, market indices are also 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*[6] 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 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[7] to explicitly identify all of the 'options' that are available. For example, in Example 3 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 and an option to increase the rate at which gold can be extracted.

Example 3 **(Luenberger's[8] 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 $\$200$ 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

[6] 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.

[7] This is another reason why there has been too much 'hype' surrounding the term "real options".

[8] This example covers Examples 12.7 and 12.10 in Luenberger's *Investment Science*. An Excel worksheet for this example, *Luenberger_Real_Options*, is in the workbook *Equity_Options.xls* that may be downloaded from the course website.

price fluctuations are described by a binomial model with each period corresponding to 1 year. Interest rates are flat at 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 given [9] on the previous page.

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 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 12.6 million and this is obtained[10] 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. 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:

[9] It is easy to create more general recombining lattices for the gold price fluctuations if necessary. For example, if it was felt to be extremely unlikely that the price of gold would ever go above $\$2,000$, then we could force the binomial lattice to respect this by only allowing very small up-moves when the gold price got sufficiently high. We could also use multinomial trees if appropriate. (Similar remarks apply to most of the examples in this course that use the binomial model.)

[10] Modulo rounding errors.

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 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. We continue working backwards in this manner, determining at each node whether or not the new equipment should be installed if it hasn't been installed already. The new lattice 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 3 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 3 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 part of the problem, i.e. the binomial lattice for gold prices, would still be complete and so we would still have unique risk-neutral (i.e. EMM) probabilities for the up-moves and down-moves in the price of gold. (We would have to use other economic considerations to help us determine[11] 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 assume that the financial market in your model is 'complete', i.e., it is *financially complete*. In this context, financial completeness now means 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.

Exercise 2 *Compute the value of the enhancement option of Example 3 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.*

The Problem Context Guides the Solution Technique

In Example 3 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[12] pricing methods available. They include[13]:

1. **Zero-Level Pricing**: Zero-level pricing[14] is based on utility maximization for portfolio optimization problems. It is that security[15] 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. 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 and are therefore not particularly suited to real options problems where hedging is usually not the primary focus. In fact these methods are more theoretical in nature and are often of more interest to the mathematical finance community. In particular, they are not suitable for the more complex real options problems that often arise in practice.

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.

[11] 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.

[12] Actually there are infinitely many EMM's available when markets are incomplete.

[13] See Chapter 16 of Luenberger for a discussion of zero-level pricing and an introduction of utility indifference pricing.

[14] Equilibrium models generally use zero-level pricing to price derivative securities in incomplete models when these derivatives are assumed to be in zero net supply in the economy.

[15] In our context, the real option is the security.

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 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[16], 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[17] probability distribution to do so and discount by the risk-free interest rate. A CAPM-based argument can provide[18] some intuition for this observation. In particular, the CAPM states

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

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[19] 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.

Remark 2 *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 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 spot oil market. Only non-financial uncertainty 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.*

---

[16] 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.

[17] If there is only private uncertainty, then it can be shown that the financial markets are irrelevant when we use zero-level pricing.

[18] As could other more sophisticated economic models.

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

Example 4 **(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

$$X_0 \quad = \quad \mathrm{E}_0^Q\left[\frac{X_1(1+r_0^{(f)})}{1+r_0^{(d)}}\right] \tag{6}$$

where X_i is the $/Euro exchange rate at time i, and $r_i^{(d)}$ and $r_i^{(f)}$ are the domestic and foreign interest rates, respectively, for the period $[i, i+1]$. Solving (6) for the risk-neutral probabilities q and $1-q$, we see that

$$q \;=\; \frac{R_p - d}{u - d} \tag{7}$$

where $R_p := (1+r_0^{(d)})/(1+r_0^{(f)})$. This of course is our usual expression for q in the binomial model, except that we now have a new pseudo-interest rate, R_p. 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 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 $\$\mu X_t$. These cash-flows are valued as usual by backwards iteration using the risk-neutral probabilities computed in (7). ∎

Exercise 3 *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?*

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 3 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.

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.

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

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

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[20] 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.

Challenge Question

Consider a broker who accepts bets on the possible outcomes of sports events. The broker is smart, successful and owns his own business. Suppose now that the broker quotes *odds* for a certain big sporting event that involves a very popular local team. The team will either win or lose and a tie is not possible. The broker has decided to quote *fair odds* so that the probabilities for winning and losing implied by his odds sum exactly to one. Given this information, do you think it is reasonable to expect that the quoted odds reflect the brokers true view regarding the probabilities of the local team winning and losing?

While interesting in its own right, this question also has implications for financial markets.

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