Chapter 7

Stochastic Discount Factors

In this chapter, we discuss ways to use information from financial markets to calibrate models for discounting future risks. This type of information is important for modeling the impact of future uncertainty on present decisions. In some areas of activity, there exist well developed financial markets with hordes of traders using the tools and information available to them to decide the present value of future events. This chapter describes a methodology to use market information.

7.1 Financial Market Information

One major question in modeling is where do discount rates come from? Financial market information consists of the prices quoted and volumes traded of financial securities—stocks, bonds, forwards, futures, and options—on public exchanges.

Securities are contracts with standard terms and conditions. For example, a share of a company is a contract that gives the owner the rights to a share of the dividends, and a bond specifies a schedule of coupon payments. Many standardized contracts used by industries are traded on exchanges.

The manufacturer of fashion goods from Chap. 6, for example, will need to ship its goods from where they are manufactured to where they will be sold. Because of the seasonality and the multiple suppliers, it is quite likely that there will be competition for transportation services in these markets. Transportation costs can be quite significant, and the unit prices can be quite variable. However, these services are entirely standard and there exist well-developed markets in shipping contracts.

It is possible to purchase *forward contracts* or *futures contracts* to ship containers of a certain dimension between major hubs like Hong Kong and New York. On the other hand, the manufacturer could choose to wait until the last minute and purchase transportation services on the *spot market*.

A.J. King and S.W. Wallace, *Modeling with Stochastic Programming*, Springer Series in ORFE, DOI 10.1007/978-0-387-87817-1_7,

Of course, relationships must exist between these prices. For example, the Hong Kong to New York price must be related to the Hong Kong to Los Angeles and the Los Angeles to New York prices. If there are any discrepancies, then speculators would enter the market and try to make money from the price differences. One of the key mechanisms by which speculators move in and out of these markets is through options.

7.1.1 A Simple Options Pricing Example

To get started, let us consider a very simple example. Our manufacturer wishes to enter an options contract to purchase a standard contract for transportation services at a future date.

To keep things simple, let us design a simple option that gives the manufacturer the right to ship 100 containers in one year's time at a fixed strike price (a call option) that we suppose to be $8,000. And suppose the manufacturer is experienced enough to know that, historically, there are really only two interesting alternatives for the spot price:

- The spot can either go up to $10,000 (because the ships are full).
- Or it can go down to $7,000 (because the ships are empty).

Now in the first case, where the spot increases in value to $10,000, the manufacturer naturally will exercise the options and save himself $2,000 in shipping costs (we will ignore the effect of interest rates on the amount invested in the option). So the value to him is the $2,000 he has saved. But if the spot decreases in value to $7,000, then the manufacturer will let the options expire, ship at the spot price, and save nothing. Table 7.1 illustrates the example.

Table 7.1: Payout diagram

Period 0	Period 1	
	Scenario 1	Future value = $10,000
		Option payout = $2,000
Forward value = $8,000		
Option value = ??		
	Scenario 2	Future value = $7,000
		Option payout = $0

How can the manufacturer calculate the value of the option? A trader would answer this question quite simply. She would say that the value of the option is the value of the portfolio that *replicates* the option payouts. Moreover, she would employ a little trick (more on this in a moment) and immediately declare that the value of the option, to the nearest penny, is \$666.67!

Before we explain the trick, we will first verify that \$666.67 would indeed replicate the option payouts. To simplify matters, we have already supposed that the value of the reference security—the bond—is always \$1. Here is how the trader does it.

1. Begin with an amount of cash \$666.67.
2. Borrow 4,666.66 units of the reference security (the bond), for a total of \$5,333.33 in cash, and purchase 2/3 forward contracts at \$8,000 per contract.
3. Liquidate the portfolio to generate the payments:
   - If the spot goes up to \$10,000 (scenario 1), then sell the 2/3 forward contract, earning \$6,666.67, and use \$4,666.66 to pay back the loan. The remaining amount in the portfolio is \$2,000.01, which is the required \$2,000 option payment (plus a penny rounding error).
   - If the spot goes down to \$7,000 (scenario 2), then liquidate the portfolio, as above. The value of the forwards sold exactly equals what you need to pay back the loan, so the amount remaining in the portfolio is \$0.

On this very simple sample space, we see that the trader can *replicate* the option payouts provided she has at least \$666.67 to start with. If the trader starts out with more than this, then under every possible scenario she will make a profit. This kind of riskless profit making is called *arbitrage*. Market models assume that there are so many traders trading that these kinds of arbitrages are eliminated by competition.

Mathematically, what we need to do is as follows: we must verify that there exists a portfolio $[x_B, x_F]$, where x_B is the number of bonds the trader buys (or borrows) and x_F the number of forwards she purchased (or sold) that satisfies the following constraint system:

$$\begin{aligned} x_B + 8{,}000x_F &= 2{,}000/3, & \\ x_B + 10{,}000x_F &\geq 2{,}000 & \text{(scenario 1)}, \\ x_B + 7{,}000x_F &\geq 0 & \text{(scenario 2)}. \end{aligned} \tag{7.1}$$

The first constraint is the budget constraint for the trader. (We have used exact numbers to make the system easier to solve.) You can verify that a solution to this system is

$$\begin{aligned} x_B &= -14{,}000/3, \\ x_F &= 2/3. \end{aligned}$$

A budget of \$666.67, and the capability to borrow a large sum of money from a bank, is what is required for the trader to replicate the option. The option cannot be replicated for less, and if we assume the market in options is free of arbitrage, then the option cannot be worth more. So one can say that the arbitrage-free value of the option is \$666.67.

But what was the trader's mental trick? Is the trader so smart that she can solve a linear program in the blink of an eye?

Well, we said there is a trick in here somewhere. Let us look again at problem (7.1). Notice that if the trader is operating in a competitive, arbitrage-free market, then she can only expect to achieve the *minimal* quantity of starting capital that replicates the option. Let us model this requirement as a linear program with an objective function that consists of this starting capital, so the goal is to use as little cash as possible to satisfy the two scenario constraints. First, express the primal constraint system in matrix form:

$$\begin{bmatrix} 1 & 10{,}000 \\ 1 & 7{,}000 \end{bmatrix} \begin{bmatrix} x_B \\ x_F \end{bmatrix} \geq \begin{bmatrix} 2{,}000 \\ 0 \end{bmatrix}. \tag{7.2}$$

Now let us develop the dual problem. (When an operations research textbook mentions a trick, you will usually find it in the dual!) Assign dual variables q^i to each of these constraints.

By linear programming duality, the dual variables are nonnegative because the primal constraint rows have lower bounds. The constraint matrix is the transpose of the primal matrix, with the right-hand side equal to the primal objective function coefficients. The dual constraint system is

$$\begin{bmatrix} 1 & 1 \\ 10{,}000 & 7{,}000 \end{bmatrix} \begin{bmatrix} q^1 \\ q^2 \end{bmatrix} = \begin{bmatrix} 1 \\ 8{,}000 \end{bmatrix}. \tag{7.3}$$

The dual constraint relationships are of an equality type because the primal variables are unconstrained, and its objective function will be a maximization with objective function coefficients equal to the right-hand side of the primal problem. It follows that the dual problem is

$$\begin{aligned} &\max_{q \geq 0} \quad 2{,}000q^1 + 0q^2 \\ &\text{such that} \quad \begin{cases} q^1 + q^2 & = 1, \\ 10{,}000q^1 + 7{,}000q^2 & = 8{,}000, \\ q^1, q^2 & \geq 0. \end{cases} \end{aligned} \tag{7.4}$$

Now let us interpret it. The first and last constraints in the dual problem show that the q's are the weights of a probability distribution because they are nonnegative and sum to one. The second condition is a kind of integration constraint that says that the expected value of the future spot prices using the probability distribution q equals the price today. The technical term for a distribution that satisfies a condition like this is *martingale*.

So this is the trader's trick! She knows this basic duality (it is quite general) and immediately sets about calculating the martingale distribution: she finds probability weights that satisfy

$$10{,}000q^1 + 7{,}000q^2 = 8{,}000.$$

There is *exactly one* answer: $[\frac{1}{3}, \frac{2}{3}]$. For this answer the value of the dual objective is

$$2{,}000\frac{1}{3} + 0\frac{2}{3} = 2{,}000/3.$$

Since there is only one answer, we can conclude that this is in fact the optimal value of the dual problem!

This is why the trader can come up with the price so quickly. She knows that in this system there is only one dual solution and that it is a martingale distribution. All she has to do is find it and compute the integral of the option payouts to determine the optimal value.

This simple options pricing trick holds a number of important points to keep in mind as the material gets more complicated.

7.1.1.1 Duality and the Martingale Distribution

First (and perhaps most) of all, options pricing is a great example of the importance of *duality* in modeling. The dual form of the options pricing problem is much easier to solve than the primal, and moreover its solution can be interpreted as a martingale distribution. In many probability settings, the martingale condition turns out to have strong implications. For instance, when asset returns are driven by exponential Brownian motions, as is the case in many applications of financial mathematics, the martingale distribution turns out to be unique. The major challenge in these settings is how to calculate the payouts and perform the integrations for the hypothesized stochastic process. But the general outline of the solution is the same as in our simple linear programming setup.

7.1.1.2 Calibration

Secondly, the example shows that prices of derivative securities in a market model are strongly related to the properties of the asset price change distribution. The martingale distribution in the simple example depends on the arrangement of future states of price changes. Of course, we chose these particular values to make a nice story. But how do traders price options in a real market? The detailed answer is complicated, of course, because in real financial markets, the details matter. The timing of dividend or interest payments, the number of trading days until expiration, carrying costs, opportunity costs, the modeling of option exercises, and the actual asset price processes all have important bearing on valuations. However, the basic answer is always the same. Options are priced in real markets by first finding the martingale distribution that best fits the observed market data. This process of inference is called *calibration*, which we will discuss in subsequent sections.

7.1.2 Stochastic Discount Factors

The pattern used by traders to price options relies on the duality between volumes of trades needed to hedge future outflows and the price distribution to apply to evaluate future outflows. The general framework of duality links the technology of trading (e.g., how efficient is it?) with structural aspects of the price distribution (e.g., is it a martingale?). In our simple example we saw that a trading program with zero transaction costs had dual solutions that were martingale distributions. Additional modeling conditions, such as transaction costs, limits on borrowing, and so forth, create a trading program with dual solutions that may no longer have the martingale property.

These results are well known and fundamental. These dual prices were first observed in their generality by Harrison and Pliska [19] in their examination of the theoretical underpinnings of market equilibrium in contingent claims. The dual variables attached to the replication conditions for the various states of the market are often called *Arrow–Debreu prices* in the economics literature. In the finance literature, the operation of finding the martingale measure and solving the integration problem is called *risk-neutral pricing* because the option is completely replicated and there is no risk. If there is risk, then it would be more appropriate to apply some form of expected utility—such as the Markowitz mean variance or the Kelly growth criterion—to the gap between what can be achieved by trading and the actual option payouts.

We prefer to use a different term that has also appeared in the literature. Our reasons are that, on the one hand, the names Arrow and Debreu would imply an economywide pricing infrastructure, and on the other hand, we do not suppose that risk vanishes in our formulations. Rather, we prefer to use the term *stochastic discount factors* (SDFs) because we intend to apply them in contexts where discounting is required and uncertainty is important.

7.1.3 Generalizing the Options Pricing Model

Let, as usual, $s \in S$ denote the scenarios, possibly developed as in Chap. 4. We let J denote the set of securities, with $j = 0$ the reference security, typically a bond. Further, the vector c denotes the present value of all the securities, while the vector c^s denotes the values in scenario s. The required value payoff from the portfolio in scenario s is f^s.

Traders price options by *replication*, in other words, they try to find the cheapest trading strategy that would replicate those option payouts. A two-stage model of option replication can be modeled as a linear program:

$$F := \min\{cx | c^s x \geq f^s, \quad s \in S\}. \tag{7.5}$$

The way this system works, from the trader's point of view, is this.

1. Begin with an amount of cash F.
2. Allocate it at today's prices, c, into a portfolio x of positions in each underlying security. Positions can be positive (long) or negative (short—a loan, in other words).
3. At each node of the scenario tree liquidate the portfolio, earning $c^s x$ to generate at least the required payment f^s.

The optimal value F is the smallest amount that can replicate all payments without ever falling short. The analysis of this simple system will enable us to derive powerful conclusions about price relationships in a market model.

SDFs are used in market models to compute the fair value of uncertain cash flows that are dependent on market observables. In this section we develop a methodology to calibrate the distribution of an SDF from information contained in market prices of related securities.

To apply the SDF methodology, we return to the two-stage replication problem (7.5) and write down its dual formulation:

$$\begin{aligned} \max_{y \geq 0} \quad & \sum_{s \in S} f^s y^s \\ \text{such that} \quad & \begin{cases} \sum_{s \in S} c_j^s y^s = c_j & j \in J \\ y^s \geq 0 & s \in S. \end{cases} \end{aligned} \tag{7.6}$$

So the dual solution $y = (y^1, \ldots, y^{|S|})$ is nonnegative and has the same dimensionality as the sample space. We call it an SDF for reasons that will become clearer in the following paragraph.

Readers familiar with the financial treatment of options pricing will recall that it is possible to recover a probability distribution from the SDF by discounting. It is not hard to see that the weights q^s given by

$$q^s := y^s \frac{c_0^s}{c_0}$$

are the weights of a probability distribution since this transforms the constraint for the reference security $j = 0$ in (7.6) into $\sum_{s \in S} q^s = 1$. We can think of an SDF in this setting as the product of two terms

$$y^s := q^s \frac{c_0}{c_0^s},$$

where q^s is a probability and c_0/c_0^s is a discount factor. The name *stochastic discount factor* seems to fit this situation well; moreover, it can also be applied to models where the dual variable is not a *martingale distribution*. We prefer to use the term SDF for all these settings.

7.1.4 Calibration of a Stochastic Discount Factor

If we can observe the market prices of related securities, then a simple way to set the weights is to express the observed prices as constraints on the SDF itself. We can write down constraints that require the integral of the option payouts to equal (or almost equal) the observed option's price, as in [36].

Suppose the market has a set I of listed options with observed bid prices, G_i^b, observed ask prices, g_i^a, and future payouts, g_i^s for $i \in I$, $s \in S$. The unknown SDF for problem (7.6) should satisfy, in addition, the following "calibration" inequalities:

$$\begin{aligned} \sum_{s\in S} y^s g_i^s &\le g_i^a, \quad i \in I, \\ \sum_{s\in S} y^s g_i^s \ge g_i^b \Leftrightarrow -\sum_{s\in S} y^s g_i^s &\le -g_i^b, \quad i \in I. \end{aligned} \tag{7.7}$$

These constraints are to be added to the constraint system of the dual problem (7.6). Each of these added constraint rows will correspond to new primal variables in the trader's replication problem: $\xi^a \ge 0$, interpreted as long positions in the listed option, and $\xi^b \ge 0$, interpreted as short positions in the listed options. Adding these inequalities results in a new primal problem. The primal objective is formed by taking the product of the dual right-hand side with the new primal variables, resulting in the following calibrated primal problem:

$$\begin{aligned} &\min_{x,\xi^a,\xi^b} \quad cx + (g^a \cdot \xi^a - g^b \cdot \xi^b) \\ &\text{such that} \quad \begin{cases} c^s x + g^s \cdot (\xi^a - \xi^b) \ge f^s & s \in S, \\ \xi_i^a, \xi_i^b \ge 0. \end{cases} \end{aligned} \tag{7.8}$$

This problem corresponds to allowing the trader to take positions in the market-traded options to hedge the payouts f^s. You may have noted that the dual constraints (7.7) are "hard." What if these constraints cannot be satisfied?

7.1.4.1 Liquidity-Weighted Calibration

It may surprise you—with all of the emphasis on arbitrage and efficient markets and equilibrium—to learn that when actual market data are entered into this problem, the solution (almost always) tries to take an unbounded position in one of the listed options. So, in fact, there is arbitrage in financial markets. How can this be?

The reason has to do with trading volumes, or market liquidity. Sellers and buyers of listed options always indicate how many option contracts they are willing to buy or sell at the advertised price. These offered volumes can be extremely low for options that are lightly traded. If anyone shows interest in buying or selling at the advertised price, then the trader will quickly pay attention. Prices for additional quantities could change substantially.

Thus, it is natural to constrain the number of options available to the model to be less than the total volume of the option traded thus far in the session.

$$\xi_i^a, \xi_i^b \leq v_i, \quad i \in I. \tag{7.9}$$

With these constraints, the solution is bounded and quite reasonable results are observed in practice. These constraints are added to the primal problem, so let us see what happens to the dual. The dual will have an additional $2I$ nonnegative variables corresponding to the liquidity bounds (7.9). Let us call these variables λ_i^a and λ_i^b for $i \in I$. Glancing back at (7.7) we can verify, through linear programming duality, that these volume bounds translate into volume-weighted penalty terms for violating the calibration inequalities, expressing the natural observation that the greater the volume traded, the more reliable are the price quotes.

$$\begin{aligned} \max_{y \geq 0} \quad & \sum_{s \in S} f^s y^s - v \cdot [\lambda^a + \lambda^b] \\ \text{such that} \quad & \begin{cases} \sum_{s \in S} c_j^s y^s & = c_j, \quad j \in J, \\ \sum_{s \in S} g_i^s y^s - \lambda_i^a & \leq g_i^a, \quad i \in I, \\ \sum_{s \in S} g_i^s y^s + \lambda_i^b & \geq g_i^b, \quad i \in I, \\ y^s & \geq 0, \quad s \in S, \\ \lambda_i^a, \lambda_i^b & \geq 0 \quad i \in I. \end{cases} \end{aligned} \tag{7.10}$$

In summary, calibration to externally observed option prices is a way of using the information in the market equilibrium to place constraints on the SDF. Calibrated discount factors can be extended by smoothing or some other statistical techniques for use as an approximate discounting operator. In the case of highly liquid markets, the SDF can be quite accurately determined from the observed market prices.

In real situations, there are two issues that must be confronted. First, the manufacturer's actual needs may not be precisely covered by the contracts being offered in the market. For example, the manufacturer may need to ship on September 1, but the contracts specify August 15 and September 15. Or the contracts may be for quantities that are different than the standard contracts. The stochastic programming approach outlined above is easily applied to these kinds of problems. A second, and more challenging, concern is that the standardized contracts may be correlated with the uncertainty but not identical with it.

7.2 Application to the Classical NewsVendor Problem

In this section, we consider how the framework for news vendor models can be extended to address uncertainty from both market and nonmarket sources. Our approach here can be viewed as an extension of Birge [6]. We utilize the model formulation of Gaur and Seshadri [17] as a starting point for this discussion.

We define the following problem variables and parameters:

$$\begin{aligned} I &: \quad \text{initial inventory to be ordered,} \\ D &: \quad \text{future demand,} \\ p &: \quad \text{unit selling price,} \\ c &: \quad \text{unit cost,} \\ s &: \quad \text{salvage price,} \\ r &: \quad \text{risk-free interest rate.} \end{aligned} \tag{7.11}$$

The news vendor has a time 0 cash flow that consists of the cost paid for the initial inventory of newspapers

$$F_0(I) = -cI \tag{7.12}$$

and a time T cash flow that consists of the income from the sold newspapers plus a salvage term for the unsold inventory:

$$F_T(I) = p\ \min[D, I] + s \max[0, I - D]. \tag{7.13}$$

At all other intervening times, $t \in (0, T)$, there are no news vendor cash flows, so we set $F_t(I) := 0$.

We suppose that a simple statistical model relates future demand D to the price at time T of a certain market-traded security, S:

$$D = bS_T + \epsilon, \tag{7.14}$$

where ϵ is a random noise term that is independent of S_T. For example, the market-traded security could be a stock-market index, which the news vendor could relate to sales: high growth of the index would indicate high sales, and poor performance could be indicative of poor sales.

We suppose that $\mathcal{M}_t$ describes the observable states for market prices S_t, and $\mathcal{N}_t$ describes the observable states for the nonmarket sources ϵ. Such a model (7.14) could be constructed by regression of demand data against the security prices. Substituting for D in (7.13) and collecting the terms yields

$$F_T(I) = pI - (p-s)(I-\epsilon) + (p-s)b\big(S_T - \max[\,0, S_T - (I-\epsilon)/b\,]\big). \quad (7.15)$$

This equation describes the news vendor's income in terms of $(p-s)b$ units of a long position in security S and $(p-s)b$ units of a short position in a "call option" with strike $(I-\epsilon)/b$. Performing some algebra we can obtain a concise expression for the news vendor's income:

$$F_T(I) = pI - (p-s)b\big[(I-\epsilon)/b - S_T + \max[\,0, S_T - (I-\epsilon)/b\,]\big] \quad (7.16)$$

$$F_T(I) = pI - (p-s)b\max[\,0, (I-\epsilon)/b - S_T\,]. \quad (7.17)$$

This shows that the news vendor time T cash flow equals the nominal income pI from selling all the inventory minus $(p-s)b$ units with cash flows that resemble the payoffs of "put options" with strike price $(I-\epsilon)/b$. This suggests that if the news vendor wanted to hedge the uncertainty in his business, he should buy put options in the market-traded security.

Put options are well-known vehicles for insuring portfolios. So what we have shown here is that the news vendor, the manufacturer desiring transportation services, or really any business that purchases inventory in advance of sales will be in the market for *insurance* in the form of put options.

Now let us suppose that the news vendor wants to sell his business. How much would someone pay for it? A buyer looks at the future cash flows and tries to determine the *most valuable* portfolio of assets he could buy in the market today that will be paid for by the cash flows from the business. (This is the reverse of the options pricing problem in which the hedger determines the minimum cost portfolio whose cash flows will exceed the cash flows of the option.) The idea here is that the buyer's future cash flows will pay for the initial asset purchase. One can express the news vendor "valuation" problem as follows:

$$\max_{I,\theta,\xi} \quad S_0\theta_0 + (C_a\xi_0^a - C_b\xi_0^b)$$
$$\text{such that} \quad \begin{cases} S_t(\theta_t - \theta_{t-1}) + (C_t\xi_0^a - C_t\xi_0^b) = F_t(I), \\ S_T\theta_T \geq 0, \\ \xi_0^a,\ \xi_0^b \geq 0, \end{cases} \quad (7.18)$$

where I^* denotes the optimal initial inventory. The optimal value V^* is the upper limit of what anyone would pay for the business.

7.2.1 Calibration of Real Options Models

This approach can also be applied to any real options modeling problem. The basic ideas are as follows:

1. Generate a scenario tree for the underlying process, for example, oil spot prices.
2. Identify some market-traded options in the underlying and model their payouts on the scenario tree.
3. Develop an investment model that allows trading in the underlying, in a bond, and allows the taking of buy-and-hold positions in the options.
4. Apply reasonable volume constraints on the options.
5. Model the cash flows of the real options problem as the underlying option to be hedged.
6. Solve for the minimum cash needed to payout the cash flows. This is the sell price.
7. Solve for the maximum cash that could be obtained from receiving the cash flows. This is the buy price.

If the sell and buy prices are close in value, then you may conclude that the model is giving you good information about the option value. If they are wide apart, then you may wish to consider incorporating a utility function into the problem. For details on this approach, see King [35] and King et al. [37].

7.3 Summary Discussion

In this section, we examine the options pricing model in the spirit of this book. After all, this is a model that claims to offer a methodology to manage uncertainty. Does it really work as advertised?

Of course, discounting is important. One can hardly plan for the future without considering discounting. Ever since the development of options pricing it has been known that the pricing of options is a natural model in which to introduce SDFs as a mechanism for accounting for the impact of future events in the consideration of the costs and benefits of current actions. But let us in any case give the model a more detailed examination.

First, what is the uncertainty? In the simple model of Sect. 7.1.1, the uncertainty space is the shipping costs:

- Forward value of spot price $8,000.
- Spot can either go up to $10,000 (because the ships are full).
- Or it can go down to $7,000 (because the ships are empty).

over which we derive the probability weights of the martingale measure:

$$\max \quad 2,000q^1 + 0q^2$$
$$\text{such that} \quad \begin{cases} q^1 + q^2 & = 1, \\ 10,000q^1 + 7,000q^2 & = 8{,}000, \\ q^1, q^2 & \geq 0, \end{cases}$$

conclude that the optimal value is \$666.67, and derive the optimal hedge: borrow \$4,666.67 and buy 2/3 of a forward contract for \$5,333.33.

Now let us look at the end stage. Denote by ξ the spot price for shipping. Here is what happens:

- The cash flow generated by the option is $\max\{0, \xi - 8{,}000\}$.
- The cash flow generated by the portfolio is $\frac{2}{3}\xi - 4{,}666.67$.

Do you see that the portfolio cash flow is linear? We have graphed the two curves in Fig. 7.1.

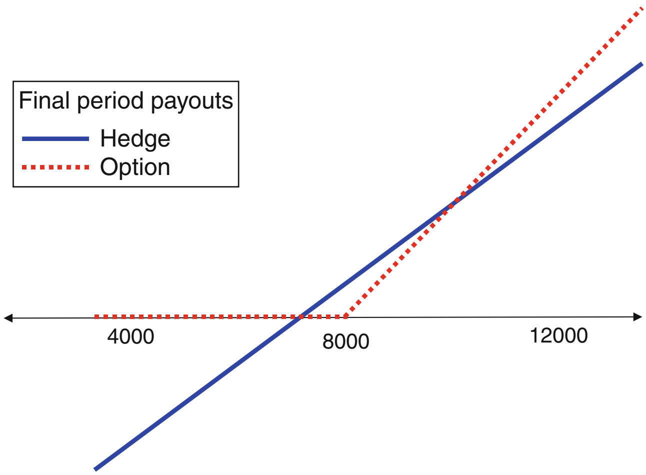

Fig. 7.1: Graph of option cash flow versus hedging portfolio cash flow

Can you see that the linear portfolio cash flow just touches the graph of the option payout at the two points corresponding to \$7,000 and \$10,000? This is the result of insisting on a hedging policy that matches cash flows at these two outcomes. But for values below \$7,000 and above \$10,000 the value of the portfolio lies below the value of the option. Our quick-minded trader did not use a model that took these possibilities into consideration!

The calibration model, on the other hand, can potentially do a better job of matching the option cash flows. For one thing, the optimization is carried out over a wider range of possible outcomes for ξ. For another, the calibration model allows the trader to take positions in listed options. These listed options have nonlinear graphs and so offer a greater capability to match the nonlinear shape of the option cash flows.

Of course, the calibration model is better precisely because some other traders have done the hard work of posting quotes for the listed options. The calibration model specifies a martingale measure using the *already known* prices of the listed options. Our quick-minded trader is *estimating* the martingale measure from a simple two-point approximation. The calibration model is making use of a great deal more information than was available to our quick-minded trader.