A System for Bank Portfolio Planning

14

Commercial banks and, to a lesser degree, other financial institutions have substantial holdings of various types of federal, state, and local government bonds. At the beginning of 1974, approximately twenty-five percent of the assets of commercial banks were held in these types of securities. Banks hold bonds for a variety of reasons. Basically, bonds provide banks with a liquidity buffer against fluctuations in demand for funds in the rest of the bank, generate needed taxable income, satisfy certain legal requirements tied to specific types of deposits, and make up a substantial part of the bank's investments that are low-risk in the eyes of the bank examiners.

In this chapter, we present a stochastic programming model to aid the investment-portfolio manager in his planning. The model does not focus on the day-to-day operational decisions of bond trading but rather on the strategic and tactical questions underlying a successful management policy over time. In the hierarchical framework presented in Chapter 5, the model is generally used for tactical planning, with certain of its constraints specified outside the model by general bank policy; the output of the model then provides guidelines for the operational aspects of daily bond trading.

The model presented here is a large-scale linear program under uncertainty. The solution procedure employs the decomposition approach presented in Chapter 12, while the solution of the resulting subproblems can be carried out by dynamicprogramming, as developed in Chapter 11. The presentation does not require knowledge of stochastic programming in general but illustrates one particular aspect of this discipline, that of ''scenario planning.'' The model is tested by managing a hypothetical portfolio of municipal bonds within the environment of historical interest rates.

14.1 OVERVIEW OF PORTFOLIO PLANNING

The bond-portfolio management problem can be viewed as a multiperiod decision problem under uncertainty, in which portfolio decisions are periodically reviewed and revised. At each decision point, the portfolio manager has an inventory of securities and funds on hand. Based on present credit-market conditions and his assessment of future interest-rate movements and demand for funds, the manager must decide which bonds to hold in the portfolio over the next time period, which bonds to sell, and which bonds to purchase from the marketplace. These decisions are made subject to constraints on total portfolio size, exposure to risk in the sense of realized and unrealized capital losses,* and other policy limitations on the makeup of the portfolio. At the next decision point, the portfolio manager faces a new set of interest rates and bond prices, and possibly new levels for the constraints, and he must then make another set of portfolio decisions that take the new information into account.

* *Realized* capital losses refer to actual losses incurred on bonds sold, while *unrealized* capital losses refer to losses that would be incurred if bonds currently held had to be sold.

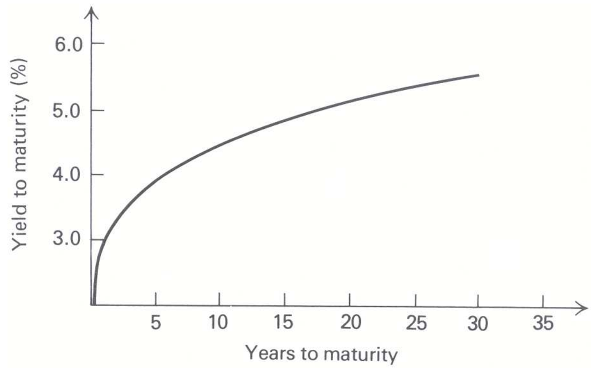

Figure 14.1 (Typical yield curve for good-grade municipal bonds.

Before describing the details of the portfolio-planning problem, it is useful to point out some of the properties of bonds. A bond is a security with a known fixed life, called its maturity, and known fixed payment schedule, usually a semiannual coupon rate plus cash value at maturity. Bonds are bought and sold in the market-place, sometimes above their face value, or par value, and sometimes below this value. If we think of a bond as having a current price, coupon schedule, and cash value at maturity, there is an internal rate of return that makes the price equal to the present value of the subsequent cash flows, including both the interest income from the coupon schedule and the cash value at maturity. This rate of return is known as the ‘‘yield to maturity’’ of a bond.

Given the attributes of a bond, knowing the price of a bond is equivalent to knowing the yield to maturity of that bond. Since the payment schedule is fixed when the bond is first issued, as bond prices rise the yield to maturity falls, and as bond prices fall the yield to maturity rises. Bond prices are a function of general market conditions and thus rise and fall with the tightening and easing of credit. Usually the fluctuations in bond prices are described in terms of yields to maturity, since these can be thought of as interest rates in the economy. Hence, bond prices are often presented in the form of yield curves. Figure 14.1 gives a typical yield curve for ‘‘good-grade’’ municipal bonds. Usually the yield curve for a particular class of securities rises with increasing maturity, reflecting higher perceived market risk associated with the longer maturities.

One final point concerns the transaction costs associated with bond trading. Bonds are purchased at the ‘‘asked’’ price and, if held to maturity, have no transaction cost. However, if bonds are sold before their maturity, they are sold at the ‘‘bid’’ price, which is lower than the ‘‘asked’’ price. The spread between these prices can be thought of as the *transaction cost* paid at the time the securities are sold.

At the heart of the portfolio-planning problem is the question of what distribution of maturities to hold during the next period and over the planning horizon in general. The difficulty of managing an investment portfolio stems not only from the uncertainty in future interest-rate movements but from the conflicting uses made of the portfolio. On the one hand, the portfolio is used to generate income, which argues for investing in the highest-yielding securities. On the other hand, the portfolio acts as a liquidity buffer, providing or absorbing funds for the rest of the bank, depending upon other demand for funds. Since this demand on the portfolio is often high when interest rates are high, a conflict occurs, since this is exactly when bond prices are low and the selling of securities could produce capital losses that a bank is unwilling to take. Since potential capital losses on longer maturities are generally higher than on shorter maturities, this argues for investing in relatively shorter maturities.

Even without using the portfolio as a liquidity buffer, there is a conflict over what distribution of maturities to hold. When interest rates are low, the bank often has a need for additional income from the portfolio; this fact argues for investing in longer maturities with their correspondingly higher yields. However, since interest rates are generally cyclical, if interest rates are expected to rise, the investment in longer maturities could build up substantial capital losses in the future, thus arguing for investing in shorter maturities. The opposite is also true. When interest rates are high, the short-term rates approach (and sometimes exceed) the long-term rates; this fact argues for investing in shorter maturities with their associated lower risk. However, if interest rates are expected to fall, this is exactly the time to invest in longer maturities with their potential for substantial capital gains in a period of falling interest rates.

Many commercial banks manage their investment portfolio using a ‘‘laddered’’ maturity structure, in which the amount invested in each maturity is the same for all maturities up to some appropriate length, say 15 years. Generally, the longer the ladder, the more risky the portfolio is considered. Figure 14.2(a) illustrates a 15-year ladder. Each year one fifteenth (i.e., $6\frac{2}{3}$ percent) of the portfolio matures and needs to be reinvested, along with the usual interest income. In a laddered portfolio, the cash from maturing securities is reinvested in fifteen-year bonds while the interest income is reinvested equally in all maturities to maintain the laddered structure. The advantages of a laddered portfolio are: no transaction costs or realized losses, since bonds are always held to maturity rather than sold; generally high interest income, since the yield curve is usually rising with increasing maturity; and ease of implementation, since theoretically no forecasting is needed and a relatively small percentage of the portfolio needs to be reinvested each year.

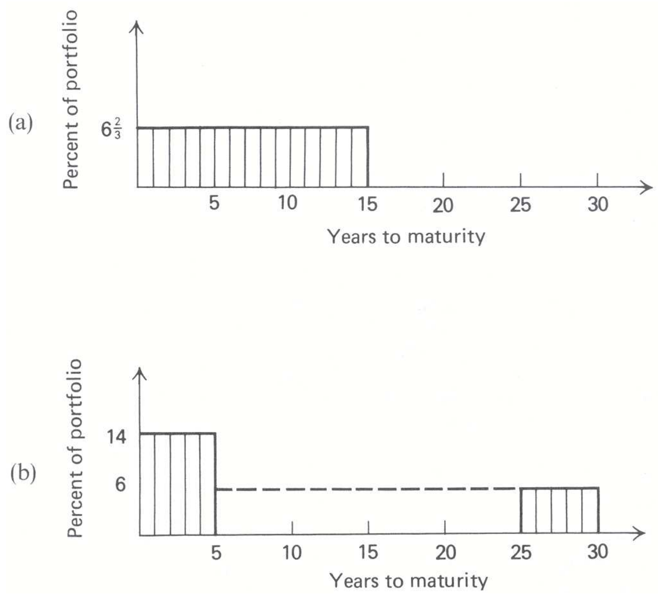

Figure 14.2 (a) Typical laddered portfolio. (b) Typical barbell portfolio.

Some banks, on the other hand, manage their portfolio using a ‘‘barbell’’ maturity structure, in which the maturities held are clustered at the short and long ends of the maturity spectrum, say 1 to 5 years and 26 to 30 years, with little if any investment in intermediate maturities. Figure 14.2(b) illustrates a typical barbell portfolio structure with 70 percent short- and 30 percent long-term maturities. The riskiness of the portfolio is judged by the percentage of the portfolio that is invested in the long maturities. Each end of the barbell portfolio is managed similarly to a ladder. On the short end, the maturing securities are reinvested in 5-year bonds, while on the long end the 25-year securities are sold and the proceeds reinvested in 30-year securities. The interest income is then used to keep the percentages of the portfolio in each maturity roughly unchanged.

The advantages of a barbell portfolio are usually stated in terms of being more ‘‘efficient’’ than a laddered portfolio. The securities on the long end provide relatively high interest income, as well as potential for capital gains in the event of falling interest rates, while the securities on the short end provide liquid assets to meet various demands for cash from the portfolio for other bank needs. In the barbell portfolio illustrated in Fig. 14.2(b), 20 percent of the portfolio is reinvested each year, since 14 percent matures on the short end and roughly 6 percent is sold on the long end. Comparing this with the $6\frac{2}{3}$ percent maturing in the 15-year ladder, it is argued that a barbell portfolio is more flexible than a laddered portfolio for meeting liquidity needs or anticipating movements in interest rates.

However, effectively managing a barbell portfolio over time presents a number of difficulties. First, significant transaction costs are associated with maintaining a barbell structure since, as time passes, the long-term securities become shorter and must be sold and reinvested in new long-term securities. Second, the short-term securities are not risk-free, since the income and capital received at maturity must be reinvested in new securities at rates that are currently uncertain. To what extent is a barbell portfolio optimal to maintain over time? One might conjecture that often it would not be advantageous to sell the long-term securities of the barbell structure and, hence, that over time the barbell would eventually evolve into a laddered structure.

In order to systematically address the question of what distribution of maturities should be held over time, a stochastic programming model was developed. The basic approach of this model, referred to as the BONDS model, is one of ''scenario planning.'' The essential idea of scenario planning is that a limited number of possible evolutions of the economy, or scenarios, is postulated, and probabilities are assigned to each. All the uncertainty in the planning process is then reduced to the question of which scenario will occur. For each scenario, a fairly complex set of attributes might have to be determined; but, given a particular scenario, these attributes are known with certainty.

We can illustrate this process by considering the tree of yield curves given in Fig. 14.3. We can define a collection of scenarios in terms of the yield curves assumed to be possible. Actually, a continuum of yield curves can occur in each of the future planning periods; however, we approximate our uncertainty as to what will occur by selecting a few representative yield curves. Suppose we say that in a three-period example, interest rates can rise, remain unchanged, or fall, in each period with equal probability. Further, although the *levels* of interest rates are serially correlated, there is satistical evidence that the distributions of *changes* in interest rates from one period to the next are independent. If we make this assumption, then there are three possible yield curves by the end of the first period, nine at the end of the second, and twenty-seven by the end of the third. (The yield curves at the end of the third period have not been shown in Fig. 14.3.) A scenario refers to *one specific sequence* of yield curves that might occur; for example, rates might rise, remain unchanged, and then fall over the three periods. The total number of scenarios in this example is $3 \times 9 \times 27$, or 729. Of course, the large number of scenarios results from our independence assumption, and it might be reasonable to eliminate some of these alternatives to reduce the problem size.

A scenario, defined by a sequence of yield curves, will have additional characteristics that place constraints on the portfolio strategy for that scenario. Since rising interest rates mean a tightening of credit, often funds are withdrawn from the portfolio, under such scenarios, to meet the demands for funds in the rest of the bank. When interest rates are falling, funds are usually plentiful, and additional funds are often made available to the portfolio. Further, a limitation on the investment strategy is imposed by the level of risk the bank is willing to tolerate. This can be expressed for each scenario by limiting the losses that may be realized within a tax year, as well as by limiting the unrealized capital losses that are allowed to build up in the portfolio over the planning horizon. Another limitation on investment strategy results from the bank's ''pledging'' requirements. The holdings of government securities, as well as the holdings of some state and local bonds, are affected by the levels of certain types of deposits. The fluctuations of these deposits are then forecast for each planning scenario, to indicate the minimum holdings of the securities that will satisfy the pledging requirements. The minimum holdings of government securities may also be affected by the bank's need for taxable income, although this taxable-income requirement also could be a characteristic of each scenario directly specified by the portfolio manager.

Scenario planning is the key to being able to concentrate on the investment portfolio. The interface between the investment portfolio and the rest of the bank is accounted for by using consistent definitions of scenarios for planning throughout the bank. For planning the investment portfolio, this interface is characterized by the demand on the portfolio for funds, the allowable levels of realized and unrealized losses in the portfolio, the limits on the holdings of certain broad categories of securities, as well as any other element of a scenario that the portfolio manager deems important for the planning problem being addressed. These characteristics of the scenarios are then tied to interest-rate movements by using the same definitions of scenarios for assessing them as for forecasting yield-curve movements. The scenario-planning process is illustrated in Section 14.4 where we discuss managing a hypothetical portfolio.

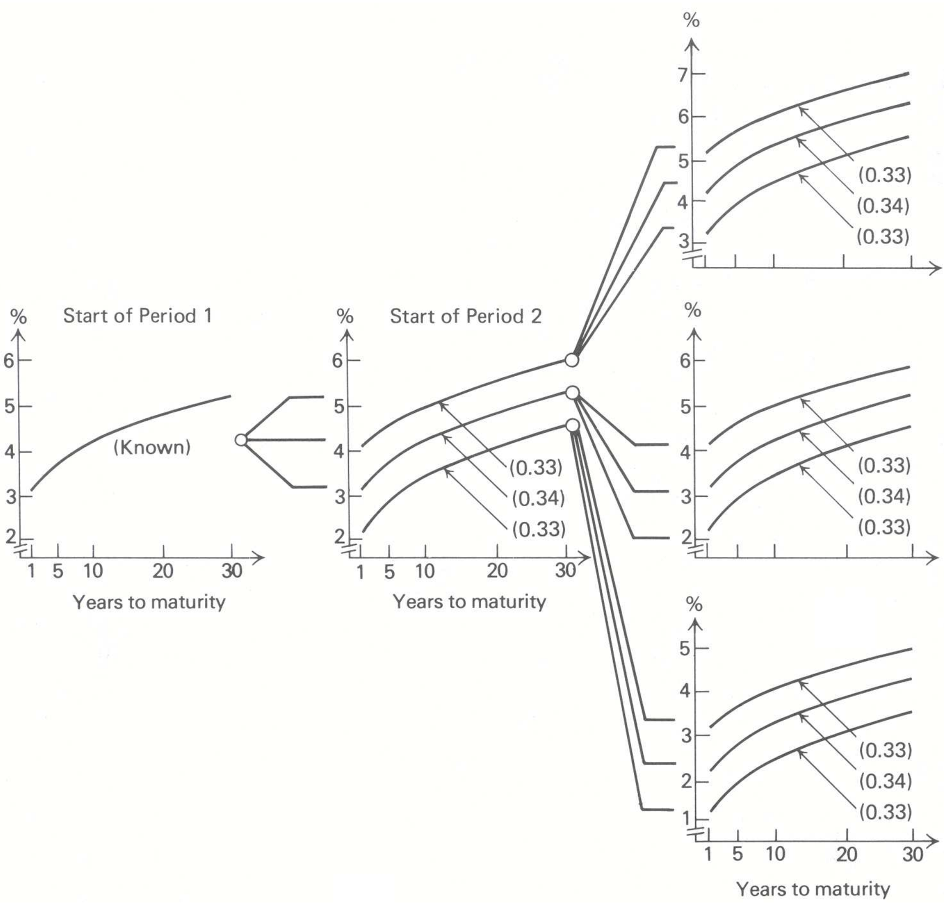

Figure 14.3 Tree of yield curves; probability in parentheses.

14.2 FORMULATION OF THE BONDS MODEL

The most important assumption in the formulation of the model is that the planning is being carried out with a limited number of economic scenarios. The scenarios are usually keyed to the movement of some appropriate short-term interest rate, such as the 90-day treasury bill rate. The possible movements of the short-term rate generate a collection of scenarios each of which consists of a particular sequence of yield curves and exogenous cash flows, as well as other characteristics for each period in the planning horizon. The assumption of a finite number of scenarios is equivalent to making a discrete approximation of the continuous distribution of *changes* in the short-term rate, and this in turn, along with the finite number of planning periods, permits the formulation of an ordinary linear program that explicitly takes uncertainty into account. Associated with any particular scenario is its probability of occurrence, which is used to structure the objective function of the linear program so as to maximize the expected horizon value of the portfolio.

The remaining characteristics of the economic scenarios are policy considerations involving the interface between the investment portfolio and the rest of the bank. For each tax year in the planning horizon, a maximum level of losses that may be realized is usually specified for each scenario. Further, the maximum level of unrealized losses that potentially could build up in the portfolio over the planning horizon is often specified. In the situation where more than one broad category of securities is being analyzed, either maximum or minimum levels of the holdings of a particular category might be specified. For example, a minimum level of U.S. Treasury holdings typically is specified, to cover the pledging of specific securities to secure certain types of state and municipal deposits.

For any particular analysis that the portfolio manager is considering, he must first group the securities to be included in the planning by broad categories, and then aggregate the securities available for purchase into a number of security classes within each category. The broad categories usually refer to securities described by the same yield curve, such as U.S. Treasury bonds or a particular grade of municipal bonds. The aggregation of securities within these broad categories is by *time to maturity*, such as 3 months, 6 months, 1 year, 2 years, . . . , 30 years. These security classes will usually not include all maturities that are available but some appropriate aggregation of these maturities.

The remainder of the section specifies the details of the mathematical formulation of the BONDS model. The discussion is divided into three parts: the decision variables, the constraints, and the objective function.

Decision Variables

At the beginning of each planning period, a particular portfolio of securities is currently held, and funds are either available for investment or required from the portfolio. The portfolio manager must decide how much of each security class k to *buy*, $b^k_n(e_n)$, and how much of each security class currently held to *sell* $s^k_{m,n}(e_n)$ or continue to *hold* $h^k_{m,n}(e_n)$. The subscript n identifies the current period and m indicates the period when the security class was purchased. Since the amount of capital gain or loss when a security class is sold will depend on the difference between its purchase price and sale price, the portfolio manager must keeptrack of the amount of each security class held, by its period of purchase. Further, since the model computes the optimal decisions at the beginning of every period for each scenario, the variables that represent decisions at the start of period n must be conditional on the scenario evolution e_n up to the start of period n. An example of a scenario evolution up to the start of period 3 would be ''interest rates rise in period 1 and remain unchanged in period 2.'' More precisely, the decision variables are defined as follows:

$b^k_n(e_n)$ = Amount of security class k *purchased* at the beginning of period n, conditional on scenario evolution e_n; in dollars of initial purchase price.

$s^k_{m,n}(e_n)$ = Amount of security class k, which had been purchased at the beginning of period m, *sold* at the beginning of period n, conditional on scenario evolution e_n; in dollars of initial purchase price.

$h^k_{m,n}(e_n)$ = Amount of security class k, which had been purchased at the beginning of period m, *held* (as opposed to sold) at the beginning of period n, conditional on scenario evolution e_n; in dollars of initial purchase price.

It should be pointed out that liabilities, as well as assets, can be included in the model at the discretion of the planner. Banks regularly borrow funds by participating in various markets open to them, such as the CD (negotiable certificate of deposit) or Eurodollar markets. The portfolio manager can then usethese ''purchased funds'' for either financing a withdrawal of funds from the portfolio or increasing the size of the portfolio. However, since the use of these funds is usually a policy decision external to the investment portfolio, an elaborate collection of liabilities is not needed. The portfolio planner may include in the model a short-term liability available in each period with maturity equal to the length of that period and cost somewhat above the price of a short-term asset with the same maturity.

Constraints

The model maximizes the expected value of the portfolio at the end of the planning horizon subject to five types of constraints on the decision variables as well as nonnegativity of thesevariables. The types of constraints, each of which will be discussed below, include the following: funds flow, inventory balance, current holdings, net capital loss (realized andunrealized), and broad category limits. In general, there are separate constraints for every time period in each of the planning scenarios. The mathematical formulation is given in Table 14.1, where e_n is a particular scenario evolution prior to period n and E_n is the set of *all possible scenario*

Table 14.1 Formulation of the BONDS model

Objective function	Maximize $\sum_{e_N \in E_N} p(e_N) \sum_{k=1}^{K} \left[\sum_{m=0}^{N-1} (y_m^k(e_m) + v_{m,N}^k(e_N)) h_{m,N}^k(e_N) + (y_N^k(e_N) + v_{N,N}^k(e_N)) b_N^k(e_N) \right]$
Funds flow	$\sum_{k=1}^{K} b_n^k(e_n) - \sum_{k=1}^{K} \left[\sum_{m=0}^{n-2} y_m^k(e_m) h_{m,n-1}^k(e_{n-1}) + y_{n-1}^k(e_{n-1}) b_{n-1}^k(e_{n-1}) \right] - \sum_{k=1}^{K} \sum_{m=0}^{n-1} (1 + g_{m,n}^k(e_n) s_{m,n}^k(e_n)) = f_n(e_n)$ $\forall e_n \in E_n \quad (n = 1, 2, \ldots, N)$
Inventory balance	$-h_{m,n-1}^k(e_{n-1}) + s_{m,n}^k(e_n) + h_{m,n}^k(e_n) = 0 \quad (m = 0, 1, \ldots, n-2)$ $-b_{n-1}^k(e_{n-1}) + s_{n-1,n}^k(e_n) + h_{n-1,n}^k(e_n) = 0$ $\forall e_n \in E_n \quad (n = 1, 2, \ldots, N; k = 1, 2, \ldots, K)$
Initial holdings	$h_{0,0}^k(e_0) = h_0^k \quad (k = 1, 2, \ldots, K)$
Capital losses	$-\sum_{k=1}^{K} \sum_{m=n'}^{n} g_{m,n}^k(e_n) s_{m,n}^k(e_n) \le L_n(e_n) \quad \forall e_n \in E_n, \quad \forall n \in N'$
Category limits	$\sum_{k \in K^i} \left[b_n^k(e_n) + \sum_{m=0}^{n-1} h_{m,n}^k(e_n) \right] \underset{(\le)}{\ge} C_n^i(e_n) \quad \forall e_n \in E_n$ $(n = 1, 2, \ldots, N; i = 1, 2, \ldots, I)$
Nonnegativity	$b_n^k(e_n) \ge 0, \quad s_{m,n}^k(e_n) \ge 0, \quad h_{m,n}^k(e_n) \ge 0 \quad \forall e_n \in E_n,$ $(m = 1, 2, \ldots, n-1; n = 1, 2, \ldots, N; k = 1, 2, \ldots, K)$

evolutions prior to period n.

Funds Flow

The funds-flow constraints require that the funds used for purchasing securities be equal to the sum of the funds generated from the coupon income on holdings during the previous period, funds generated from sales of securities, and exogenous funds flow. We need to assess coefficients reflecting the income *yield* stemming from the semiannual coupon interest from holding a security and the capital *gain or loss* from selling a security, where each is expressed as a percent of initial purchase price. It is assumed that taxes are paid when income and/or gains are received, so that these coefficients are defined as after-tax. Transaction costs are taken into account by adjusting the gain coefficient for the broker's commission; i.e., bonds are purchased at the ''asked'' price and sold at the ''bid'' price. We also need to assess the exogenous funds flow, reflecting changes in the level of funds made available to the portfolio. The exogenous funds flow may be either positive or negative, depending on whether funds are being made available to or withdrawn from the portfolio, respectively.

The income yield from coupon interest, the capital gain or loss from selling a security, and the exogenous funds flow can be defined as follows:

$g_{m,n}^k(e_n)$ = Capital gain or loss on security class k, which had been purchased at the beginning of period m and was sold at the beginning of period n conditional on scenario evolution e_n; per dollar of initial purchase price.

$y_m^k(e_n)$ = Income yield from interest coupons on security class k, which was purchased at the beginning of period m, conditional on scenario evolution e_n; per dollar of initial purchase price.

$f_n(e_n)$ = Incremental amount of funds either made available to or withdrawn from the portfolio at the beginning of period n, conditional on scenario evolution e_n; in dollars.

Since it is always possible to purchase a one-period security that has no transaction cost, the funds-flow constraints hold with equality implying that the portfolio is at all times fully invested. Finally, if short-term liabilities are included in the model, then these funds-flow constraints would also reflect the possibility of generating additional funds by selling a one-period liability.

Inventory Balance

The current holdings of each security class purchased in a particular period need to be accounted for in order to compute capital gains and losses. The inventory-balance constraints state that the amount of these holdings sold, plus the remaining amount held at the beginning of a period, must equal the amount on hand at the end of the previous period. The amount on hand at the end of the previous period is either the amount purchased at the beginning of the previous period or the amount held from an earlier purchase.

It is important to point out that this formulation of the problem includes security classes that mature before the time horizon of the model. This is accomplished by setting the hold variable for a matured security to zero (actually dropping the variable from the model). This has the effect, through the inventory-balance constraints, of forcing the ‘‘sale’’ of the security at the time the security matures. In this case, the gain coefficient reflects the fact that there are no transaction costs when securities mature.

Initial Holdings

The inventory-balance constraints also allow us to take into account the securities held in the initial portfolio. If the amounts of these holdings are:

h_0^k = Amount of security class k held in the initial portfolio; in dollars of initial purchase price,

the values of the variables that refer to the holdings of securities in the initial portfolio, $h_{0,0}^k(e_0)$, are set to these amounts.

Capital Losses

Theoretically, we might like to maximize the bank's expected utility for coupon income and capital gains over time. However, such a function would be difficult for a portfolio manager to specify; and further, management would be unlikely to have much confidence in recommendations based on such a theoretical construct. Therefore, in lieu of management's utility function, a set of constraints is added that limit the net *realized* capital loss during any year, as well as the net *unrealized* capital loss that is allowed to build up over the planning horizon.

Loss constraints are particularly appropriate for banks, in part because of a general aversion to capital losses, but also because of capital adequacy and tax considerations. Measures of adequate bank capital, such as that of the Federal Reserve Board of Governors, relate the amount of capital required to the amount of ‘‘risk’’ in the bank's assets. Thus, a bank's capital position affects its willingness to hold assets with capital-loss potential. Further, capital losses can be offset against taxable income to reduce the size of the after-tax loss by roughly 50 percent. As a result, the amount of taxable income, which is sometimes relatively small in commercial banks, imposes an upper limit on the level of losses a bank is willing to absorb.

The loss constraints sum over the periods contained in a particular year the gains or losses from sales of securities in that year, and limit this value to:

$L_n(e_n)$ = Upper bound on the realized net capital loss (after taxes) from sales during the year ending with period n, conditional on scenario evolution e_n; in dollars.

In Table 14.1, N' is the set of indices of periods that correspond to the end of fiscal years, and n' is the index of the first period in a year defined by an element of N'. Thus the loss constraints sum the losses incurred in all periods that make up a fiscal year. Since the model forces the ''sale'' of all securities at the horizon without transaction costs, the unrealized loss constraints have the same form as the realized loss constraints.

Category Limits

It may be of considerable interest to segment the portfolio into a number of broad asset categories each of which is described by different yield curves, transaction costs, income-tax rates, and distribution of maturities. There is no conceptual difficulty with this; however, some computational difficulties may arise due to problem size. Typically the investment portfolio might be segmented into U.S. Treasury and tax-exempt securities. In addition, the tax-exempt securities might be further segmented by quality such as prime-, good-, and medium-grade municipals. In making such a segmentation, we often impose upper- or lower-bound constraints on the total holdings of some of the asset categories. The example cited earlier involved placing lower limits on the amount of U.S. Treasury bonds held for pledging purposes. Defining

$C_n^i(e_i)$ = Lower (upper) bound on the level of holdings of asset category i, at the beginning of period n, conditional on scenario evolution e_n; in dollars of initial purchase price,

and letting K^i be the index set of the ith asset category, the constraints in Table ?? merely place an upper or lower bound on the holdings of a particular broad asset category.

Objective Function

The objective of the model is to maximize the expected value of the portfolio at the end of the final period. It should be pointed out that this assumes that the portfolio manager is indifferent between revenues received from interest and revenues received from capital gains, since each add equivalent dollars to the value at the horizon. If desired, it would be possible to include interest-income constraints to ensure that sufficient income would be achieved by the portfolio during each period.

The final value of the portfolio consists of the interest income received in the final period, plus the value of securities held at the horizon. It is not obvious, though, how the value of the portfolio holdings should be measured, since they are likely to contain some unrealized gains and losses. Should these gains or losses be calculated before or after taxes? Before or after transaction costs? At one extreme, it would be possible to assume that the portfolio would be sold at the horizon, so that its final value would be after taxes and transaction costs. This approach would tend to artifically encourage short maturities in the portfolio, since they have low transaction costs. The alternative approach of valuing the portfolio before taxes and transaction costs in equally unrealistic. For simplicity, it is usually assumed that the value of the securities at the horizon is after taxes but before transaction costs.

The objective function can be defined in terms of the decisions made at the start of the final period, which are conditional on the evolution of each scenario up to that point. For each scenario, the value of any holdings at the start of the final period should reflect the expected capital gain or loss over the final period and the coupon income to be received in the final period. The objective function can be formalized by defining the *probability* of scenario evolution and the *value* of the noncash holdings at the start of the final period as follows:

$v_{m,N}(e_N)$ = Expected (over period N) cash value perdollar of initial purchase price of security class k, which had been purchased at the beginning of period m, and held at the start of period N, conditional on scenario evolution e_n;

$p(e_N)$ = Probability that scenario evolution e_N occurs prior to period N.

The expected horizon value of the portfolio given in Table 14.1 is then determined by weight-ing the value of the holdings at the start of the final period by the probability of each scenario.

14.3 PROBLEM SIZE AND STRUCTURE

In order to have a feeling for the potential size of the model formulated, the number of constraints implied under various model assumptions can be computed. Assume for the moment that the number of events in each time period is the same, and equal to D. Thus there are D scenario evolutions for the first period, D^2 for the second, and so forth. Further, let there be a total of n time periods with n_i periods in year i. Then if K is the total number of different security classes in all categories, and I is the number of broad asset categories, the number of equations can be calculated as follows:

Cash flow:	$1 + D + D^2 + \cdots + D^{n-1}$
Net capital loss:	$D^{n_1} + D^{n_1+n_2} + \cdots$
Category limits:	$(I-1)\left[1 + D + D^2 + \cdots + D^{n-1}\right]$
Inventory balance:	$K[D + 2D^2 + \cdots + nD^{n-1}]$

Table 14.2 indicates the number of each type of constraint under a variety of assumptions. It is clear that, for even a relatively small number of events and time periods, the problem size rapidly becomes completely unmanageable. However, it is also clear that the main difficulty lies with the number of inventory-balance constraints. Hence, an efficient solution procedure is likely to treat these constraints implicitly instead of explicitly.

Table 14.2 Number of constraints in various models

	$I = 2, \quad K = 8$					
	$D = 3$			$D = 5$		
Period/year	$n_1 = 1,\ n_2 = 1$ $n_3 = 1$	$n_1 = 2,\ n_2 = 1$ $n_3 = 1$	$n_1 = 3,\ n_2 = 1$ $n_3 = 1$	$n_1 = 2,\ n_2 = 1$ $n_3 = 1$	$n_1 = 2,\ n_2 = 1$ $n_3 = 1$	$n_1 = 3,\ n_2 = 1$ $n_3 = 1$
Cash flow	13	40	121	31	160	781
Net capital loss	39	117	350	155	775	3,875
Category limit	13	40	121	31	160	781
Inventory balance	168	816	3408	440	3440	23,440
Total constraints	233	1613	4000	657	4535	28,877

Figure 14.4 illustrates the structure of the linear-programming tableau for a model involving three time periods and three securities. Note that the inventory-balance constraints exhibit a block diagonal structure. Given this structure, the inventory-balance constraints can be treated implicitly rather than explicitly by a number of techniques of large-scale mathematical programming. For this particular application, the decom-position approach, which was introduced in Chapter 12, was used, since the resulting special structure of the subproblems could be exploited readily, both in the solution of the subproblems and in structuring the restricted master. In this approach, there is one subproblem for each diagonal block of inventory-balance

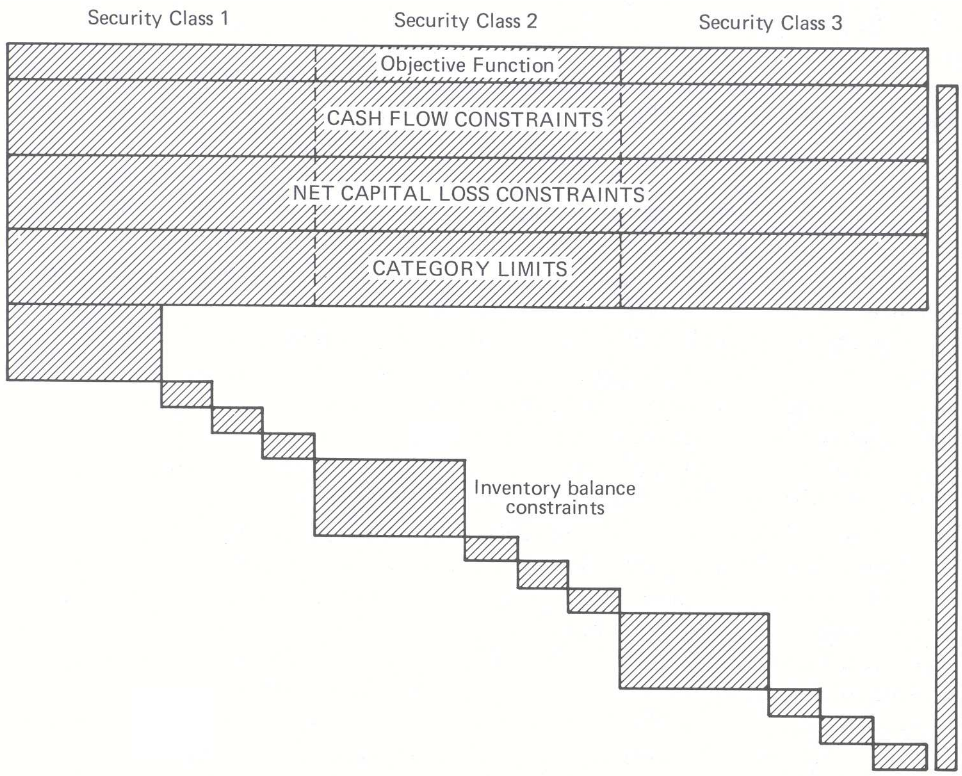

Figure 14.4 Tableau structure for three security classes.

constraints, while the restricted master linear program that must be solved at each iteration has constraints corresponding to the funds flow, net capital loss, and category limit constraints.

The subproblems correspond to purchasing each security class at the start of a given period, conditional on every scenario evolution up to that period. For example, purchasing security class k at the start of period one defines one subproblem and purchasing the same security class at the start of period two defines one additional subproblem for each scenario evolution in period one. This is indicated in the tableau structure given in Fig. 14.4 by the three small rectangles following one large one on the diagonal. The decision variables of a subproblem involve selling and holding in subsequent periods the security class purchased. This allows capital gains and losses on sales of securities to be determined.

To illustrate the subproblems further, consider first the period 1 subproblems. Security class k is available for purchase at the start of period 1. If a decision to purchase is made, this security class is then available for sale at the start of period 2, or it may be held during the second period. The amount that is held is then available for sale or holding at the start of the third period. This multistage problem, involving a sequence of sell and hold decisions after a purchase, can be solved by a recursive procedure. The problem has a dynamic-programming structure (see Chapter 11) where the state of the systemat time n is defined by the amount of the initial purchase still held. This amount is constrained by the inventory-balance equations, which limit the amount sold in any period to be less than or equal to the amount on hand.

Note that if security class k is purchased at the start of period 2, its purchase price and income yield are conditional on the scenario evolution which occurred during period 1. Thus, one subproblem is defined for each security class that can be purchased and every possible scenario evolution that precedes the purchase period of that class. As would be expected, the subproblems have no decision variables in common with one another, since each set of inventory-balance constraints simply keeps track of the remaining holdings of a particular purchase. This approach leads to a relatively large number of subproblems; however, the rationale is that the subproblems should be efficient to solve, since the state variable of each is one-dimensional.

Another interesting point to note about the subproblem constraints is that they are homogeneous systems of equations (i.e., zero righthand sides). As we saw in Chapter 12, the fundamental theorem employed in decomposition is that the feasible region defined by a system of linear equations may be represented by a convex combination of its extreme points, plus a nonnegative combination of its extreme rays. The solutions of any subproblem have only one extreme point, all decision variables equal to zero. For any nonzero point satisfying the subproblem constraints, a scalar times that point also satisfies the constraints; and hence, with a linear objective function, there exists an associated unbounded solution. As a result we need consider only the extreme rays of the subproblems. These extreme rays may be constructed in an efficient manner either by dynamic programming or by exploiting the triangular structure of the dual of a related ‘‘ray-finding’’ linear program. The ray-finding problem is defined by setting the ‘‘buy’’ variable of any subproblem to one, and then determining the optimal sequence of sell and hold decisions for this one unit.

The restricted master for this decomposition scheme reflects the fact that only extreme rays of the subproblems need to be considered. The usual constraints that require that the solution be a convex combination of extreme points of the subproblems are not necessary, since the only restriction on the weights on the extreme rays is that they be nonnegative. When the model is solved, all profitable rays found at an iteration are added as columns to the restricted master. The restricted master is then solved by continuing the simplex method from the previous solution, which yields new shadow prices, or dual variables. The shadow prices are then used to modify the objective functions of the subproblems and the process is repeated. If no profitable ray is found for any subproblem, then the algorithm terminates, and we have an optimal solution.

The value of the optimal solution is merely given by the nonnegative weighted combination of the subproblem solutions when the weights are determined by the values of the variables in the final restricted master. In general, we need *not* add any unprofitable ray to the restricted master. However, a ray that is unprofitable at one iteration may become profitable at a future iteration, as the objective functions of the subproblems are modified by the shadow prices. Hence, if one profitable ray is generated for any subproblem at an iteration, all new rays generated, profitable or not, are in fact added to the restricted master as columns.

As the restricted master is augmented by more and more columns, those columns not in the current basis are retained, provided storage limitations permit. As storage limitations become binding, those columns that price out most negatively are dropped. Any dropped column will be regenerated automatically if needed.

Further details in the computational procedure are included in the exercises.

14.4 MANAGING A HYPOTHETICAL PORTFOLIO

In the remainder of this chapter, the use of the BONDS model is illustrated by addressing the question of what portfolio strategy should be adopted over time. The model is used to ‘‘manage’’ a hypothetical portfolio of municipal securities over a 10-year historical period. The important advantage derived from using such a model to aid in planning the portfolio maturity structure is that it gives the portfolio manager the opportunity to take explicitly into account the characteristics of the current portfolio, as well as expected interest-rate swings, liquidity needs, programs for realized losses, and exposure to unrealized losses.

Ideally, the performance of the model should be evaluated using Monte Carlo simulation. However, such an experiment would involve a significant number of simulation trials, where portfolio revisions would have to be made by the optimization model at the beginning of each year of the simulation. Updating and running the BONDS model such a large number of times would be prohibitively expensive, from the standpoints of both computer and analyst time.

As an alternative to the Monte Carlo simulation, it is possible to perform a *historical* simulation, which considers how interest rates actually behaved over a particular period of time,and then attempts to plan a portfolio strategy that could have been followed over this period. To implement this approach, the ten-year historical period starting with January 1, 1964 was chosen. In order to keep the simulation simple, portfolio decisions were allowed to be made once a year and the resulting portfolio was held for the entire year. It is not suggested that any bank would have followed the strategy proposed by the model for the entire year; however, it can be considered a rough approximation of such a strategy over the ten-year period.

It should be strongly emphasized that it is difficult to draw firm conclusions from the results of a historical

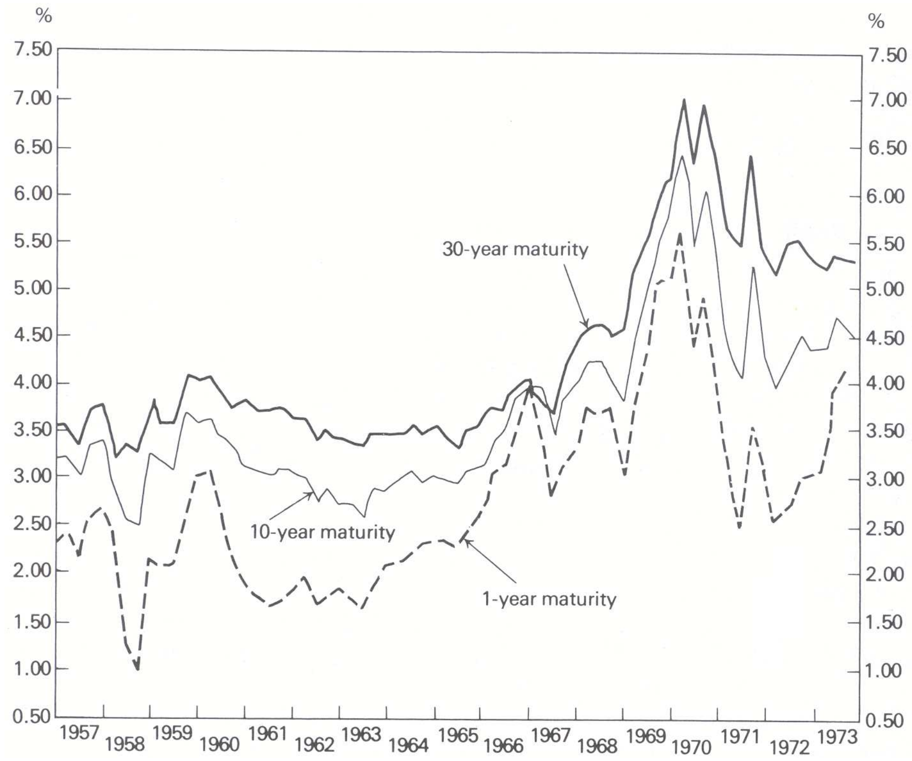

Figure 14.5 Yields of good-grade municipals.

simulation against one particular realization of interest rates. A strategy that performed well against that particular sequence of rates might have performed poorly against some other sequence of rates that had a relatively high likelihood of occurring. The opposite is, of course, also true. However, it does allow us to make comparisons between strategies for a particular sequence of interest rates that indeed did occur.

The historical simulation covered January 1964 through January 1974, a period of ten years. However, a number of years of history prior to the beginning of the simulation period were included, since it was necessary to assess interest-rate expectations for a portfolio manager at the beginning of 1964. Figure 14.5 gives the yields on one-, ten-, and thirty-year maturities for good-grade municipal bonds covering the appropriate period. Although some years exhibited a great deal of variation within a year, with a potential for improving performance through an active trading policy, this variation was not included, since in the historical simulation the portfolio was revised only at the beginning of each year.

The basic approach of the historical simulation was to use the BONDS model to make portfolio decisions at the beginning of each year, given the current actual yield curve and the portfolio manager's ''reasonable expectations'' of future interest-rate movements. These decisions were then implemented, the actual performance of the portfolio in that year was revealed, and the process was repeated. There were two steps in modeling the yield curves needed for the simulation. First, eleven actual yield curves were of interest—one for the beginning of each year when portfolio decisions were made, and one for the final performance evaluation. These yield curves were developed by fitting a functional form to historical data. Second, the portfolio manager's reasonable expectations about future interest-rate fluctuations were modeled by constructing a tree of yield curves, similar to those given in Fig. 14.3, for each year in the simulation.

Modeling the eleven yield curves that occurred was relatively straightforward. Data were available from the actual yield curves at each point in time covering the 1-, 2-, 5-, 10-, 20-, and 30-year maturities. Since yield curves are generally considered to be smooth curves, the data for each curve were fitted with the following functional form:

$$R_m = am^b e^{cm},$$

where R_m is the yield to maturity on securities with m years to maturity, and a, b, and c are constants to be determined from the data. In the simulation, the yield for any maturity was then taken from the derived yield curve for the appropriate year.

Modeling the tree of yield curves reflecting the portfolio manager's reasonable expectations of future interest-rate movements was more complicated. In the historical simulation, portfolio decisions were made as if it were January 1964; it was important not to use any information that was *not available* to a portfolio manager at that time. It is difficult to imagine or reconstruct exactly what a portfolio manager would have forecast for future interest-rate changes in January 1964. Therefore, for the purpose of the simulation, the portfolio manager's interest-rate assessments were mechanically based on the previous seven years of data at each stage. The simulation started with the interest-rate data for the years 1957 through and including 1963; and, based on these data, a tree of yield curves reflecting the portfolio manager's expectations of interest rates was constructed. This tree of yield curves was then used in making portfolio decisions for the year 1964. For each subsequent year, the data corresponding to the year just past was added to the data base and the data more than seven years old were dropped; a new tree of yield curves was then constructed.

Two separate analyses were performed to estimate the tree of yield curves representing the portfolio manager's reasonable expectations of interest rates at the beginning of each of the ten years of the simulation.

First, the distributions of one-year changes in the one-year rate were estimated for each year in the simulation. A monthly time series covering the prior seven years was used to determine the actual distribution of the one-year changes in the one-year rate. Table 14.3 gives the means and standard deviations of these distributions.

Table 14.3 Distribution of one-year changes in one-year rate

Year	*Mean (Basis points)**	*Standard deviation (Basis points)*
′64	−0.2	74.6
′65	−3.3	72.4
′66	12.4	59.8
′67	15.2	61.9
′68	11.6	59.0
′69	24.4	52.3
′70	41.9	68.6
′71	37.7	83.5
′72	11.2	109.5
′73	6.7	110.5

*100 basis-point change equals 1 percentage-point change.

Second, the changes in two other rates, the twenty- and thirty-year rates, were forecast, *conditional* on the changes in the one-year rate. Then, given a forecast change in the one-year rate, three points on the corresponding forecast future-yield curve could be determined, by adding the current levels for these rates to the forecast changes in these rates. The new levels for the one-, twenty-, and thirty-year rates then determined the constants a, b, and c for the functional form of the yield curve given above.

To forecast the changes in the twenty- and thirty-year rates, conditional on the changes in the one-year rate, two regressions were performed for each year in the simulation. The changes in each of these two rates were separately regressed against changes in the one-year rate. The two regression equations were then used to compute the changes in these rates as deterministic functionsof the changes in the one-year rate. Table 14.4 shows the means and standard deviations of the regression coefficients as well as a goodness-of-fit measure. The mean of the regression coefficient gives the change in the twenty-year rate or the change in the thirty-year rate as a fraction of the change in the one-year rate. Given a forecast change in the one-year rate, forecasts of the changes in the twenty- and thirty-year rates were determined by multiplying the change in the one-year rate by the appropriate fraction.

Table 14.4 Changes in the 20- and 30-year rates as a fraction of the changes in the one-year rate

	20-year fraction			*30-year fraction*		
Year	*Mean*	*Standard deviation*	*R-square*	*Mean*	*Standard deviation*	*R-square*
′64	.3615	.0416	85.6	.3289	.0398	86.8
′65	.2820	.0312	91.9	.2473	.0274	93.7
′66	.3010	.0354	89.1	.2762	.0298	92.3
′67	.3862	.0331	90.4	.3222	.0261	94.0
′68	.3578	.0399	86.3	.2792	.0358	88.9
′69	.4413	.0501	74.7	.3609	.0504	74.3
′70	.6753	.0421	79.8	.6187	.0478	74.0
′71	.6777	.0507	74.3	.6176	.0549	69.8
′72	.6585	.0431	84.4	.6004	.0468	81.6
′73	.6702	.0437	84.1	.6119	.0478	81.0

An important question to address regarding the assessment of future scenarios for each year of the simulation concerned the handling of the time trend in the interest-rate data. In Table14.4, the mean of the distribution of one-year changes in the one-year rate is almost always positive, indicating increasing rates, on the average. Would a portfolio manager have assumed that interest rates would continue to increase according to these historical means, when he was forecasting future rates at the beginning of each year of the simulation? It was assumed that a portfolio manager would have initially forecast the upward drift in rates prior to 1970, corresponding to the mean based on the previous seven years of data. In 1969, interest rates dropped precipitously and uncertainty increased about the direction of future changes. Hence, from 1970 until the end of the simulation, it was assumed that the portfolio manager would have forecast no net drift in interest rates, but would have had a variance of that forecast corresponding to that observed in the previous seven years.

Finally, a tree of scenarios for bank planning purposes was determined at each stage by making a discrete approximation of the distribution of changes in the one-year rate. For ease of computation, it was assumed that there were only three possible changes that could occur during each period in the planning horizon. Further, the distribution of one-year changes in the one-year rate was approximately normal; this allowed these changes to be approximated by the mean and the mean plus-or-minus one standard deviation.

By approximating the distribution of one-year changes in the one-year rate with three points, the number of branches on the scenario-planning tree at each stage is 3 in the first period, 9 in the second period, and 27 in the third period. The method just described generates a yield curve for each branch of the planning tree similar to those given in Fig. 14.3. Normally, theportfolio manager would also assess the exogenous cash flow either to or from the portfolio on each branch of the scenario. However, since the main interest was in evaluating the performance of the portfolio, the assumption was made that all cash generated by the portfolio was reinvested, and that no net cash was either made available to or withdrawn from the portfolio after the initial investment.

The only securities under consideration in the stimulation were good-grade municipal bonds. The purchase of nine different maturities was allowed—1-, 2-, 3-, 4-, 5-, 10-, 15-, 20-, and 30-years. This is a robust enough collection to show how the model behaves, although at times some of the maturities that were not included might have been slightly preferable. Trading these securities involves a cost, which is paid at the point of sale of the securities and amounts to the spread between the ‘‘bid" and ‘‘asked" prices of the market. For bond prices quoted in terms of $100 per bond, the bid–asked spread ranged from $\frac{1}{8}$ for bonds with two years or less to maturity, to $\frac{3}{4}$ for bonds with 11 years or more to maturity.

The planning horizon for the bank was taken to be three years. Since this was a yearly simulation, three one-year periods were then used in structuring the model. Assuming a three-year horizon for planning purposes indicates that the bank is willing to say that its objective is to maximize the value of the portfolio at the end of three years. Therefore, at any point in time, the bank is planning as if it wants to maximize its portfolio value at the end of the next three years, but in fact it is always rolling over the horizon, so the end is never reached.

Throughout the simulation, the limit on realized losses within any one year on any scenario was held to 0.5 percent of the initial book value of the portfolio. That is, at the beginning of each year in the simulation, the current book value of the portfolio was computed and the losses that were allowed to be realized that year and planned for over each of the years of the planning horizon were limited to 0.5 percent of this value.

The unrealized losses were handled differently. Since the actual interest rates rose over the course of the simulation, fairly large amounts of unrealized losses tended to build up in the portfolio. The potential unrealized losses were constrained to be the same for all scenarios, and the level of these unrealized losses, under moderately adverse circumstances, was kept as small as possible. With this view, the limits on potential unrealized losses were as low as 1 percent and as high as 5 percent, depending on how much the interest rates had risen to date. In the first year of the simulation, the allowable unrealized losses on all scenarios at the end of the three years in the planning horizon were limited to 1 percent of the initial book value.

14.5 EVALUATING THE PERFORMANCE OF THE MODEL

Rather than examine the details of the transactions each time the portfolio was revised, let us consider only the general structure of the portfolio over time. Figure 14.6 illustrates theholdings of the portfolio over time in broad maturity categories. It is interesting to note the extent to which a roughly barbell portfolio strategy was maintained. Initially, some intermediate maturities (15- and 20-year) were purchased. However, the 15-year maturity was sold off as soon as the program for realized losses would allow, and the 20-year maturity was gradually sold off. No other investments in intermediate securities were made during the simulation. The final portfolio was essentially a barbell structure with 1-, 2-, 3-, and 5-year maturities on the short end; and 26-, 28-, and 30-year maturities on the long end.

We can get an idea of how the value of the portfolio increased over time from Table 14.5. It should be emphasized that the period of the simulation shows a very large increase in interest rates in general, and therefore a potential for large capital losses for programs involving active trading of securities. The generally high interest income exhibited by the portfolio is not completely reflected in the increased book value of the portfolio the next year, because losses resulting from trading have reduced the book value of the portfolio. However, the year-to-year increase in the book value of the portfolio generally follows the interest-rate pattern. The final value of $141,269 for the portfolio in January 1974 corresponds to a compounded rate of return on the initial investment of 3.52 percent per year.

Table 14.5 also indicates the level of the losses realized by selling securities, and the limit placed on the level of unrealized losses that could potentially build up in the portfolio. The realized losses were constrained in each year on every scenario to be less than or equal to one-half percent of the current book value of the portfolio. In the years 1966 through 1971, the maximum level of realized losses was attained. In each year the potential unrealized losses at the horizon three years hence were approximately equal to the level of losses already built up in the portfolio.

The performance of the BONDS model can be compared with the results of applying various laddered and barbell strategies to the same historical interest-rate sequence. Alternative portfolio strategies were implemented essentially the same way as in the BONDS model, except that no forecasting procedure was needed since the strategies were applied in a mechanical manner. For the laddered portfolios, at the beginning of each year of the simulation, the funds from maturing securities were reinvested in the longest security allowed in the ladder, and the coupon income was distributed equally among all maturities to keep the

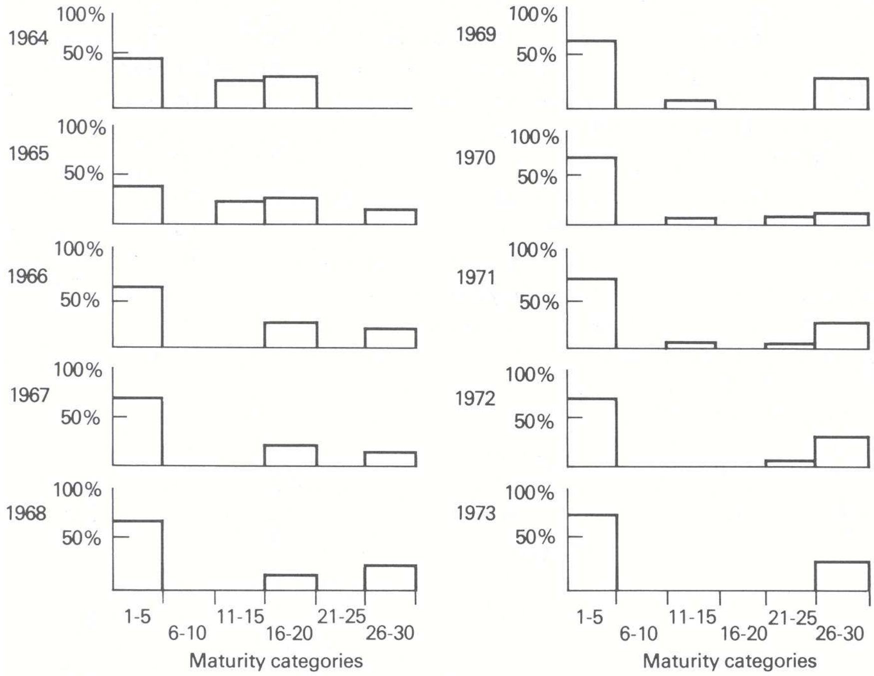

Figure 14.6 Maturity distributions of managed portfolios (beginning of year).

proportions in each fixed. For the barbell portfolios, at the beginning of each year of the simulation, the shortest securities in the long end of the barbell were sold and reinvested in the longest security allowed. Similarly, the maturing securities were reinvested in the longest security allowed on the short end. The coupon income was allocated between the long or short ends of the barbell to maintain the stated proportions, and within either end it was distributed equally among the maturities included.

Table 14.6 summarizes the results of using various laddered portfolio strategies in the same historical environment. Any relatively short ladder would seem to be a good strategy against such a rising interest-rate sequence. The five-year ladder gives the highest final value, which is slightly less than that obtained by the BONDS model. However, had the interest-rate sequence been decreasing or had it shown more fluctuation, the BONDS model would have performed significantly better than the five-year ladder, by being able to balance capital losses against gains in each period while retaining a relatively large proportion of the portfolio in longer maturities with higher interest income.

Table 14.7 summarizes the performance of using various barbell portfolio strategies in the same historical environment. The column labeled ‘‘Strategy’’ gives the maturities held in the portfolio at all times and the percentage of the portfolio that was reinvested in the long end of the barbell. For example, 1–7, 24–30, 20% means a barbell portfolio structure with seven bonds on the short end, and seven bonds on the long end, but with only twenty percent of the total value of the portfolio invested in the long end.

It is useful to compare laddered portfolios with barbell portfolios having the same number of bonds on the short end. It is clear that, for this particular realization of interest rates, the laddered portfolios did better than the comparable barbell portfolios. This is due to the fact that, although the barbell portfolios had higher interest incomes, these were offset by having to realize losses from selling a relatively long maturity during a period of increasing interest rates. If interest rates had fallen, the barbell portfolios would still have had higher average interest incomes than the ladders but would have realized capital gains rather than losses. Hence, for this historicalsimulation in a period of rising rates, the laddered portfolios outperformed the barbell portfolios but this will not generally be the case.

Within any barbell structure, as the percent of the portfolio invested in the long end was increased, the performance of the portfolio deteriorated. This is again due to the fact that the realized losses were so large

Table 14.5 Portfolio characteristics for each year

Year	*Beginning book value*	*Interest income*	*Income return on book value*	*Realized net losses*	*Potential unrealized losses at horizon*
1964	$100,000.	$2962.	2.96%	$ 0	1.0%
1965	102,959.	3150.	3.06	3	1.0
1966	105,578.	3296.	3.12	531.*	1.5
1967	108,327.	3808.	3.52	547.*	1.5
1968	111,574.	4305.	3.86	561.*	2.0
1969	115,300.	4601.	3.99	579.*	2.5
1970	119,301.	5935.	4.97	600.*	5.0
1971	124,610.	5358.	4.30	626.*	2.5
1972	129,631.	5635.	4.35	337.	2.5
1973	134,729.	6017.	4.47	537.	2.5
1974	141,269.	—	—	−523.	—

* Maximum allowable realized losses taken.

Table 14.6 Laddered portfolio strategies

Strategy	*Total interest*	*Total realized losses* (A/T)	*Final unrealized gain* (A/T)	*Final value*	*Compound rate of return*
1–10	$39,775.	$0	$ 642.	$140,416	3.45%
1–7	40,553.	0	345.	140,898	3.49
1–5	41,355.	0	−175.	141,180	3.51
1–3	41,043.	0	−409.	140,634	3.47
1	37,596.	0	0	137,596	3.24

Table 14.7 Barbell portfolio strategies

Strategy	*Total interest*	*Total realized losses* (A/T)	*Final unrealized gain* (A/T)	*Final value*	*Compound rate of return*
1–10, 21–30, 20%	$40,949.	$1723.	$ 189.	$139,415	3.38%
40	42,135.	3456.	−266.	138,413	3.30
1–7, 24–30, 20	41,879.	2579.	635.	139,935	3.42
40	43,220.	5178.	931.	138,973	3.35
1–5, 26–30, 20	42,835.	2953.	458.	140,340	3.45
40	44,334.	5933.	1102.	139,503	3.39
1–3, 28–30, 20	42,896.	2422.	−415.	140,059	3.43
40	44,775.	4865.	−420.	139,490	3.38
1, 30, 20	40,242.	3020.	0	137,222	3.21
40	42,941.	6066.	0	136,875	3.19

over this period of rising interest rates. Larger amounts invested in the long end produced larger total interest income, since the yield curves were generally increasing with longer maturities. However, the increased interest income was not sufficient to offset the capital losses incurred in trading.

Finally, the riskiness of the portfolio strategies can be analyzed by looking at the amount of unrealized losses that had built up under each strategy in the year of peak interest rates, 1970. Table 14.8 gives the book value, market value, and unrealized after-tax losses as a percent of book value, for the strategy followed by the BONDS model and for various laddered and barbell strategies. In a period of rising rates, strategies that keep most of their assets in short-term securities would be expected to have lower unrealized losses. In the extreme case, a ladder consisting of only the one-year maturity would have *zero* unrealized losses at the beginning of each year, since the entire portfolio would mature at the end of each period.

Table 14.8 Unrealized losses, January 1, 1970

Type of portfolio	*Book value*	*Market value*	*Unrealized after-tax losses*
BONDS	$119,299	$107,536	4.93%
Ladder 1–5	119,198	115,252	1.66
1–10	119,403	113,958	4.56
Barbell 1–5, 26–30, 20%	118,051	113,688	3.70
1–5, 26–30, 40%	116,889	110,107	5.80
Barbell 1–10, 21–30, 20%	119,222	111,200	6.73
1–10, 21–30, 40%	119,039	108,415	8.92

On this dimension, the strategy followed by the BONDS model builds up significant losses, roughly comparable to a ten-year ladder, or a barbell with five maturities on each end and, say, thirty percent of the portfolio value invested in the long end. However, this general level of unrealized losses in the credit crunch of 1969 might have been considered very reasonable. Any strategy that places a significant proportion of its assets in relatively short-term securities will not develop unrealized losses. However, just as there is potential for losses with these strategies, there is also potential for similar gains.

We can sum up the performance of the BONDS model by comparing it to the best of the mechanically-managed laddered or barbell strategies. The performance of the five-year laddered portfolio was almost as good as that of the BONDS model against the particular sequence of interest rates that occurred. However, it should be noted that a portfolio laddered out only to five years is a very conservative strategy, which is rather unlikely to be followed in practice. It happened to turn out, after the fact, that this was a good strategy to have adopted; but it is not clear that any portfolio manager would have been motivated to adopt such a conservative strategy consistently over the period of the simulation.

Finally, had interest rates been level or falling, the BONDS model would certainly have outperformed the five-year ladder. In these instances, some other laddered or barbell portfolio might perform competitively with the BONDS model. However, the important characteristic of the BONDS model is that it produces a strategy that is adaptive to the environment over time. The BONDS model should perform well against *any* interest-rate sequence, while a particular laddered or barbell portfolio will perform well for some realizations of interest rates and poorly for others. In actual practice, the portfolio manager would be *actively forecasting* interest rates, and the BONDS model provides a method of systematically taking advantage of these forecasts.

EXERCISES

1. A bank considering using the BONDS model to aid in portfolio planning divides its portfolio into two pools of funds—U.S. Governments and all grades of municipals. The bank managers forecast a yield curve for their group of municipals and are willing to treat them as one category. The investment portfolio decisions are revised monthly, but no securities with less than three months to maturity are purchased.

   a) Assuming that the bank employs a two-year planning horizon with four planning periods of 3 months, 3 months, 6 months, and one year, how many constraints of each type will their model have, using 5-point approximations to the uncertainties in the first two periods and 3-point approximations in the last two periods? (Assume no initial holdings, but include the remaining constraints of Table 14.1.)

b) The bank feels that it can aggregate the purchase decisions on individual securities into the following maturities: 3 months, 6 months, 1, 2, 3, 5, 10, and 20 years for U.S. Governments and 1, 2, 3, 4, 5, 10, 20, and 30 years for the municipal group. How many decision variables will the model sketched in (a) have? (Again, assume no initial holdings.)

c) In fact, the bank does have initial holdings of both Governments and municipals. The bank is willing to aggregate these holdings in the same maturities as used for new purchases. How many additional constraints and variables need to be added to account for the initial holdings?

d) How is a subproblem defined for the model described in (a), (b), and (c)? How many such subproblems are there? Why is it impossible to combine the subproblem from initial holdings with those of subsequent purchases?

e) How many constraints does the restricted master have? How would you find an initial basic feasible solution to the restricted master?

2. For the model described in Exercise 1, the demands for information from the portfolio manager are extensive.

a) How many yield curves need to be assessed to use the proposed planning model?

b) How would you think about the problem of assessing the exogenous cash flows in such a way that they are consistent with the yield curves assessed?

c) What is the interpretation of using the same level of realized and unrealized losses on all scenarios? Which scenarios are likely to produce the binding realized and unrealized loss constraints? Can the loss constraints on the remaining scenarios be dropped?

d) Suppose that the lower-bound constraint on the holdings of government securities could be dropped from the model. How might this change in the model formulation affect the optimal solution? Could the value of the objective function decrease?

3. The objective function of the BONDS model is given in terms of the horizon value of the portfolio. Suppose that we wish to reformulate the model in terms of the interest income and capital gains in each period. The decision variables remain the same; all constraints are unchanged, but the objective function changes.

a) Formulate a new objective function that maximizes the expected *increase* in the size of the portfolio by summing, over all time periods, the interest income plus capital gains (or losses) in each time period.

b) Show that the two formulations are equivalent, in the sense that the optimal values of the decision variable using each objective function are identical.

c) Show that the difference between the optimal values of the two objective functions is the expected exogenous cash flow.

4. Upon seeing your formulation in Exercise 2, the portfolio manager argues that the cash flows in the objective function should be discounted.

a) Will discounting the cash flows change the optimal values of the decision variables?

b) What are the arguments for and against discounting in this situation? How could a single discount rate be chosen?

c) What is the relationship between the shadow prices on the funds-flow constraints of the undiscounted formulation and the discount rate? (See exercises in Chapter 4.)

d) Suppose you win the argument with the portfolio manager and do not use a discounted objective function. Other departments in the bank place demands on the portfolio for funds in various time periods. Can you suggest an approach to choosing among various requests for funds? How does the approach relate to discounting the cash flows?

5. Consider the subproblems generated by the decomposition approach. Formulate the subproblem corresponding to buying a 20-year U.S. Government security at the beginning of the first period of a model consisting of 3 one-year periods. The generic decision variables to use are as follows:

$$b_1, \quad S_{21}(e_2), \quad h_{21}(e_2), \quad S_{31}(e_3), \quad h_{31}(e_3).$$

(Do not include buying a similar security at the beginning of the second or third periods.)

a) How many constraints and decision variables does the subproblem have?
b) The constraints of the subproblems are homogeneous (i.e., zero righthand sides). Suppose that purchasing 1 unit of this security, $b_1 = 1$, gives a positive rate of return. What can be said about purchasing λb_1 units of this security?
c) Formulate a dynamic-programmig model to solve this subproblem, assuming that $b_1 = 1$. Show that this solution determines a *ray* of the subproblem.

6. Suppose that, for the subproblems formulated in Exercise 5, we define a *ray-finding* subproblem as follows: b_1 is set equal to 1 and moved to the righthand side; the resulting subproblem is solved by linear programming.

a) Formulate the ray-finding problem.
b) Find the dual of the ray-finding problem.
c) Show that a basis for the dual problem is triangular.
d) Write down a recursive method for calculating the optimal solution of the dual of the ray-finding problem. [*Hint*. Exploit the triangular property of the basis to solve for the dual variables by back-substitution.]
e) How is the solution of the primal ray-finding problem determined from the solution of the dual?

7. As a follow-up to Exercises 5 and 6, once the subproblems have been solved, the rays generated are added to the restricted master:

a) Describe how the columns of the restricted master are computed.
b) Why are weighting constraints *not* needed for the restricted master in this particular application?
c) What is the stopping rule for this variation of the decomposition algorithm?
d) Suppose that storage limitations force you to drop some of the nonbasic columns from the restricted master at some iteration. Is it possible that the algorithm will be unable to find the overall optimal solution as a result?

8. There are at least two ways to define subproblems for this model. First, a subproblem can be defined for each maturity and each year it can be purchased. In Fig. 14.4 there are a total of 12 subproblems using this definition. Second, a subproblem can be defined for each maturity regardless of when it is purchased. In Fig. 14.4 there are a total of 3 subproblems using this definition.

a) Explain why the choice of subproblem definitions need have no impact on the solution procedure adopted for the subproblems.
b) Explain how the restricted master will differ under each definition of the subproblmes.
c) Which choice of subproblem definitions will make the restricted master more efficient to solve? Why?
d) If there was a weighting constraint for each subproblem, how would your answer to (c) be affected? [*Hint*. Which definition would add more weighting constraints?]

9. Suppose that a new objective function is to be considered that is a nonlinear function of the holdings entering the final period. More precisely, assume that we wish to maximize the *expected* utility of these holdings, with the utility function given by:

$$u(x) = ax^{1-c},$$

where $a > 0$ and $0 < c < 1$. The objective function is then of the form:

$$\sum_{e_N \in E_N} p(e_N) u \left\{ \sum_{k=1}^{K} \left[\sum_{m=0}^{N-1} (y_m^k(e_m) + v_{m,N}^k(e_N)) h_{m,N}^k(e_N) + (y_N^k(e_N) + v_{N,N}^k(e_N)) b_N^k(e_N) \right] \right\}.$$

a) Show that this problem cannot be handled directly by decomposition. [*Hint*. Is this objective function separable in the appropriate way?]
b) If the nonlinear-programming problem were solved by the Frank-Wolfe algorithm, a sequence of linear programs would be solved. How can the decomposition approach presented in this chapter be used to solve one of these linear programs?
c) How can the Frank-Wolfe algorithm be efficiently combined with the decomposition approach presented in this chapter, to find the optimal solution to the nonlinear program defined by maximizing the expected utility given above?

d) Does your proposed method generalize to other nonlinear problems?

10. Describe the method of producing scenarios used in the historical simulation over the first three decisions (two years).

    a) How many yield curves need to be assessed at the beginning of ’64 for the years ’64, ’65, and ’66?
    b) Draw a decision tree depicting the probability that each of the yield curves indicated in (a) occurs. What use is made of the uncertainty in the estimates of the change in the twenty-year rate (thirty-year rate) as a fraction of the change in the one-year rate? How could this be modified?
    c) Show how to compute the parameters of the functional form for each of the yield curves forecast for the start of ’65.
    d) Explain how you might assess the required information in actual practice, using a model similar to that described in Exercise 1. Draw a distinction between short-term forecasting (3 months) and long-term forecasting (2 years).

11. The performance of the laddered portfolio structure with five equally spaced maturities proved to be the best among the ladders investigated. Similarly, the performance of the 1–5, 26–30 barbell, with 20% invested in the long end, proved the best among the barbells. The BONDS model outperformed these strategies in the historical simulation over a period of rising rates.

    a) How would you expect these strategies to compare with the BONDS model in a period of falling rates?
    b) How would you choose an appropriate ladder or barbell strategy without the benefit of 20/20 hindsight?
    c) What benefits does the BONDS model have over either a ladder or barbell approach to portfolio management?

ACKNOWLEDGMENTS

This chapter is based primarily on material drawn from ‘‘Managing a Bank Bond Portfolio Over Time,’’ by S. P. Bradley and D. B. Crane, to appear in *Stochastic Programming* (ed. M. A. H. Dempster), Academic Press, in press. A broad coverage of the area is contained in *Management of Bank Portfolios* by S. P. Bradley and D. B. Crane, John Wiley & Sons, New York, 1975.