Dynamic Programming Models:

Multi-Period Portfolio Optimization

Sirong Luo

Faculty of Statistics and Data Sciences
Shanghai University of Finance and Economics

Introduction

This chapter describes four types of dynamic portfolio optimization problems that are amenable to dynamic programming technology. The first two deal respectively with optimization of final wealth and its extension, optimal consumption and investment. These two classical models date back multiple decades. The last two problems are much more modern developments. One of them is a model for dynamic trading when returns are predictable and trading is costly. The other one is a model for dynamic portfolio optimization that incorporates capital gains taxes.

Utility of Terminal Wealth

Let us revisit the dynamic portfolio optimization model with initial endowment and utility of terminal wealth that we discussed in Chapter 12. However, this time we consider a more general setting where there are forecasting variables available at each stage. Such forecasting variables could be associated with a factor model. For instance, they could be macroeconomic indicators, or certain measurable parameters of a particular asset or firm.

As before, suppose that an investor starts at $t = 0$ with an initial endowment W_0. At times $t = 0, \ldots, T-1$ the investor invests her wealth W_t in a portfolio of risk-free and risky assets. The investor's goal is to maximize the expected utility of terminal wealth $U(W_T)$ at time T for some utility function $U(\cdot)$. Define the following convenient notation:

- $R_{f,t+1}$ = gross risk-free return in period $[t, t+1]$;
- $\mathbf{r}_{t+1}$ = vector of excess returns of the risky assets in period $[t, t+1]$;
- $R_{p,t+1}$ = gross random return of the investor's portfolio in period $[t, t+1]$;
- $\mathbf{z}_t$ = forecasting state variables available at stage t;
- W_t = wealth at stage t.

We have the inter-temporal budget constraint:

$$W_{t+1} = W_t \cdot R_{p,t+1} = W_t \cdot \left(R_{f,t+1} + \mathbf{r}_{t+1}^\top \mathbf{x}_t\right), \quad t = 0, \ldots, T-1$$

The specific components of this sequential decision problem are as follows:

Stages: these are $t = 0, \ldots, T-1$.

State at stage t : this is $(W_t, \mathbf{z}_t)$.

Decision variables at stage t : these are the vector $\mathbf{x}_t$ of portfolio holdings (percentages) in the risky assets.

Law of motion: this is the same as the above inter-temporal constraint

$$W_{t+1} = W_t \cdot \left(R_{f,t+1} + \mathbf{r}_{t+1}^\top \mathbf{x}_t\right), t = 0, \ldots, T-1$$

Bellman's optimality principle

We next apply Bellman's optimality principle. In this case the value-to-go function is

$$\begin{aligned} J_t\left(W_t, \mathbf{z}_t\right) &= \max_{\mathbf{x}_t, \ldots, \mathbf{x}_{T-1}} \mathbb{E}_t\left(U\left(W_T\right)\right) \\ &= \max_{\mathbf{x}_t, \ldots, \mathbf{x}_{T-1}} \mathbb{E}_t\left[U\left(W_t \cdot \prod_{\tau=t}^{T-1}\left(R_{f,\tau+1} + \mathbf{r}_{\tau+1}^{\top}\mathbf{x}_\tau\right)\right)\right] \end{aligned}$$

At the final stage T we get

$$J_T\left(W_T, \mathbf{z}_T\right) = U\left(W_T\right).$$

For earlier stages, we have the Bellman equation

$$\begin{aligned} J_t\left(W_t, \mathbf{z}_t\right) &= \max_{\mathbf{x}_t} \mathbb{E}_t\left[J_{t+1}\left(W_{t+1}, \mathbf{z}_{t+1}\right)\right] \\ &= \max_{\mathbf{x}_t} \mathbb{E}_t\left[J_{t+1}\left(W_t\left(R_{f,t+1} + \mathbf{r}_{t+1}^{\top}\mathbf{x}_t\right), \mathbf{z}_{t+1}\right)\right] \end{aligned}$$

In the special case of power utility $U(W) = W^{1-\gamma}/(1-\gamma)$, where $\gamma > 0$, we rewrite the Bellman equation as follows. Define $\psi_t(\mathbf{z}_t) := J_t(1, \mathbf{z}_t)$. Then it is easy to see that the Bellman equation is equivalent to

$$\psi_t(\mathbf{z}_t) = \max_{\mathbf{x}_t} \mathbb{E}_t \left[\left(R_{f,t+1} + \mathbf{r}_{t+1}^\top \mathbf{x}_t \right)^{1-\gamma} \cdot \psi_{t+1}(\mathbf{z}_{t+1}) \right]$$

We can draw the following interesting conclusions from here. On the one hand, if $\mathbf{r}_{t+1}$ and $\mathbf{z}_{t+1}$ are independent at time t, then the term on the right-hand side above satisfies

$$\begin{aligned} &\mathbb{E}_t \left[\left(R_{f,t+1} + \mathbf{r}_{t+1}^\top \mathbf{x}_t \right)^{1-\gamma} \cdot \psi_{t+1}(\mathbf{z}_{t+1}) \right] \\ &\quad = \mathbb{E}_t \left[\left(R_{f,t+1} + \mathbf{r}_{t+1}^\top \mathbf{x}_t \right)^{1-\gamma} \right] \cdot \mathbb{E}_t \left(\psi_{t+1}(\mathbf{z}_{t+1}) \right) \end{aligned} \tag{14.1}$$

Thus, to find $\mathbf{x}_t$ we need to solve

$$\max_{\mathbf{x}_t} \mathbb{E}_t \left[\frac{\left(R_{f,t+1} + \mathbf{r}_{t+1}^\top \mathbf{x}_t \right)^{1-\gamma}}{1-\gamma} \right] = \max_{\mathbf{x}_t} \mathbb{E}_t \left[U(R_{p,t+1}) \right]$$

In this case the optimal policy is myopic.

On the other hand, if $\mathbf{r}_{t+1}$ and $\mathbf{z}_{t+1}$ are correlated, then (14.1) no longer holds. In this case $\mathbf{x}_t$ may include some kind of "inter-temporal hedging component". The intuition is that the correlation between $\mathbf{r}_{t+1}$ and $\mathbf{z}_{t+1}$ would induce some kind of serial dependence in our returns. In other words, the current forecasted return $\mathbf{r}_{t+1}$ conveys information about future returns. Unlike the myopic strategy, the optimal dynamic strategy incorporates this serial dependence.

Optimal Consumption and Investment

Consider an extension of the previous dynamic portfolio optimization model where the goal is to maximize an expected utility that combines two terms: consumption along the planning horizon and terminal wealth. The latter component is sometimes called bequest.

There are three key differences from the previous model. First, there is an additional decision variable $C_t \in [0, W_t]$ at each stage t that denotes the amount of wealth the investor consumes at stage t. Second, the objective function is

$$\max_{\substack{\mathbf{x}_0,\ldots,x_{T-1}\\ C_0,\ldots,C_{T-1}}} \mathbb{E}\left(\sum_{t=0}^{T-1} U\left(C_t\right) + B\left(W_T\right)\right)$$

for some utility functions $U(C)$ and $B(W)$. Third, the new law of motion, or inter-temporal budget constraint, is

$$W_{t+1} = \left(W_t - C_t\right) \cdot R_{p,t+1} = \left(W_t - C_t\right) \cdot \left(R_{f,t+1} + \mathbf{r}_{t+1}^{\top}\mathbf{x}_t\right), t = 0, \ldots, T-1$$

To simplify our discussion we consider the case when there are no forecasting variables $\mathbf{z}_t$. In particular this implies that the returns on the risky assets are independent across different time periods. At the final stage T we have the following value-to-go function

$$J_T\left(W_T\right) = B\left(W_T\right)$$

For earlier stages, we have the Bellman equation

$$J_t\left(W_t\right) = \max_{C_t, \mathbf{x}_t} \mathbb{E}_t\left[J_{t+1}\left(W_{t+1}\right) + U\left(C_t\right)\right]$$

The first-order optimality conditions yield

$$U'\left(C_t\right) = \mathbb{E}_t\left[J'_{t+1}\left(W_{t+1}\right) R_{p,t+1}\right]$$

and

$$\mathbb{E}_t\left[J'_{t+1}\left(W_{t+1}\right) \mathbf{r}_{t+1}\right] = \mathbf{0}$$

The first one is obtained by differentiating with respect to C_t and the second one is obtained by differentiating with respect to $\mathbf{x}_t$.

If we plug the optimal $C_t, \mathbf{x}_t$ back into the Bellman equation and differentiate with respect to the state variable W_t, we obtain the following envelope condition:

$$U'\left(C_t\right) = J_t'\left(W_t\right)$$

In the special case of a logarithmic utility of consumption and bequest $U(C) = \log(C), B(W) = \log(W)$, we can draw a more explicit conclusion about the problem. In this case the Bellman equation yields the following expressions for the value function and optimal consumption:

$$J_t\left(W_t\right) = \frac{\log\left(W_t\right)}{T-t+1} + b_t$$

and

$$C_t^*\left(W_t\right) = \frac{W_t}{T-t+1}.$$

The specific value b_t and the optimal portfolio $\mathbf{x}_t^*\left(W_t\right)$ depend on the joint probability distribution of $R_{f,t+1}$ and $\mathbf{r}_{t+1}$. By contrast, the optimal consumption $C_t^*\left(W_t\right)$ only depends on W_t.

Dynamic Trading with Predictable Returns and Transaction Costs

We next discuss a recent model due to Gârleanu and Pedersen (2013) for dynamic portfolio optimization when asset returns are predictable by signals and trading is costly. This problem is quite timely and especially relevant for active investors. The optimal trading policy should balance various tradeoffs. Fast trading generates more alpha and lower risk but also higher transaction costs. Slow trading does the opposite. On the other hand, there may be fast signals that require quick action and slow signals associated with longer-lasting alpha. The model that we discuss next provides an insightful solution to this problem.

Consider a universe of assets, whose returns evolve according to the following law of motion:

$$\mathbf{r}_{t+1} = \mathbf{B}\mathbf{f}_t + \mathbf{u}_{t+1}$$

Here $\mathbf{f}_t$ is a vector of factor returns that predict asset returns, $\mathbf{B}$ is a matrix of exposures or sensitivities of the asset returns to factor returns, and $\mathbf{u}_t$ is an idiosyncratic zero-mean noise term with constant covariance matrix

$$\text{var}_t\left(\mathbf{u}_{t+1}\right) := \Sigma.$$

The vector of factor returns $\mathbf{f}_t$ is known to the investor at time t and evolves according to

$$\Delta \mathbf{f}_{t+1} = -\Phi \mathbf{f}_t + \epsilon_{t+1},$$

where $\Delta \mathbf{f}_{t+1} = \mathbf{f}_{t+1} - \mathbf{f}_t$. Trading is costly. The transaction cost associated with trading the vector of shares $\Delta \mathrm{x}_t = \mathrm{x}_t - \mathrm{x}_{t-1}$ is

$$TC\left(\Delta \mathbf{x}_t\right) = \frac{1}{2} \Delta \mathbf{x}_t^\top \Lambda \Delta \mathbf{x}_t$$

for some symmetric positive definite matrix Λ. The model objective is

$$\max_{\mathbf{x}_0, \mathbf{x}_1, \ldots} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} (1-\rho)^{t+1} \left(\mathbf{r}_{t+1}^\top \mathbf{x}_t - \frac{\gamma}{2} \mathbf{x}_t^\top \Sigma \mathbf{x}_t \right) - \frac{(1-\rho)^t}{2} \Delta \mathbf{x}_t^\top \Lambda \Delta \mathbf{x}_t \right].$$

Gârleanu and Pedersen (2013) apply a dynamic programming approach to characterize the optimal trading strategy. We summarize the main results below.

The state at time t is the pair $(\mathbf{x}_{t-1}, \mathbf{f}_t)$. The value-to-go function is

$$V(\mathbf{x}_{t-1}, \mathbf{f}_t) = \max_{\mathbf{x}_t, \mathbf{x}_{t+1}, \ldots} \mathbb{E}_t \left[\sum_{\tau=t}^{\infty} (1-\rho)^{\tau+1-t} \left(\mathbf{r}_{\tau+1}^\top \mathbf{x}_\tau - \frac{\gamma}{2} \mathbf{x}_\tau^\top \Sigma \mathbf{x}_\tau \right) \right.$$
$$\left. - \frac{(1-\rho)^{\tau-t}}{2} \Delta \mathbf{x}_\tau^\top \Lambda \Delta \mathbf{x}_\tau \right]$$

Hence the Bellman equation is

$$V(\mathbf{x}_{t-1}, \mathbf{f}_t) = \max_{\mathbf{x}_t} \left\{ -\frac{1}{2} \Delta \mathbf{x}_t^\top \Lambda \Delta \mathbf{x}_t \right.$$
$$\left. + (1-\rho) \left(\mathbb{E}_t \left(\mathbf{r}_{t+1}^\top \mathbf{x}_t \right) - \frac{\gamma}{2} \mathbf{x}_t^\top \Sigma \mathbf{x}_t + \mathbb{E}_t \left[V(\mathbf{x}_t, \mathbf{f}_{t+1}) \right] \right) \right\}$$

We make an educated guess and later verify (ansatz) the following quadratic form for the value function:

$$V(\mathbf{x}_t, \mathbf{f}_{t+1}) = -\frac{1}{2} \mathbf{x}_t^\top \mathbf{A}_{xx} \mathbf{x}_t + \mathbf{x}_t^\top \mathbf{A}_{xf} \mathbf{f}_{t+1} + \frac{1}{2} \mathbf{f}_{t+1}^\top \mathbf{A}_{ff} \mathbf{f}_{t+1} + a_0$$

Using this ansatz, it can be shown that the optimal trading policy is

$$\mathbf{x}_t = \mathbf{x}_{t-1} + \Lambda^{-1}\mathbf{A}_{xx}\left(\mathrm{aim}_t - \mathbf{x}_{t-1}\right)$$

where

$$\mathrm{aim}_t = \mathbf{A}_{xx}^{-1}\mathbf{A}_{xf}\mathbf{f}_t$$

The Bellman equation also yields expressions for the matrices $\mathbf{A}_{xx}, \mathbf{A}_{xf}, \mathbf{A}_{ff}$. See the exercises at the end of the chapter. In the special case $\Lambda = \lambda\Sigma$ we obtain

$$\mathbf{x}_t = \left(1 - \frac{a}{\lambda}\right)\mathbf{x}_{t-1} + \frac{a}{\lambda}\,\mathrm{aim}_t$$

where

$$a = \frac{-(\gamma(1-\rho) + \lambda\rho) + \sqrt{(\gamma(1-\rho) + \lambda\rho)^2 + 4\gamma\lambda(1-\rho)^2}}{2(1-\rho)}.$$

Next, we get a more explicit expression of the aim portfolio. To that end, first observe that the myopic solution in the absence of transaction costs is precisely the solution to the static Markowitz model at time t; that is,

$$\text{Markowitz}_t = (\gamma\Sigma)^{-1}\mathbf{B}\mathbf{f}_t.$$

Again we consider the special case $\Lambda = \lambda\Sigma$. For $z := \gamma/(\gamma + a)$ we get

$$\begin{aligned}\text{aim}_t &= z \cdot \text{Markowitz}_t + (1-z)\mathbb{E}_t\left(\text{aim}_{t+1}\right)\\ &= \sum_{\tau=t}^{\infty} z(1-z)^{\tau-t}\mathbb{E}_t\left(\text{Markowitz}_\tau\right)\end{aligned}$$

Furthermore, the portfolio aim_t has a similar form to Markowitz_t provided the forecasting signals are appropriately scaled down:

$$\text{aim}_t = (\gamma\Sigma)^{-1}\mathbf{B}\left(\mathbf{I} + \frac{a}{\gamma}\Phi\right)^{-1}\mathbf{f}_t$$

The optimal strategy is characterized by two principles. First, aim in front of the target. Second, trade partially towards the current aim. More precisely, the optimal updated portfolio is a linear combination of the existing portfolio and an aim portfolio. The latter is a weighted average of the current Markowitz portfolio (the moving target) and the expected Markowitz portfolios on all future dates (where the target is moving).

Taxes pose a significant friction to most investors in financial markets. There are a variety of taxes that apply in different ways to income, dividends, and capital gains. It is common to ignore taxes in traditional finance and portfolio theory. This simplification is in part due to the difficulties involved in modeling the effects of taxes.

Capital taxes introduce a peculiar type of challenge in portfolio management. Since the sale of an appreciated asset triggers a capital gain tax liability, there is a tradeoff between the benefits of diversification versus the tax costs triggered by rebalancing the portfolio. In addition to the tradeoff between diversification and taxes, many individual investors also have to deal with both a tax-deferred and a taxable account. In this context an investor faces an asset location problem in addition to the usual asset allocation problem. Asset location refers to the problem of how the investor should locate her portfolio holdings across the taxdeferred and taxable accounts.

Basic Case: Tax Management Only

In the United States tax code, capital gains and losses are triggered when assets are sold. This feature means that the investor could manage her assets in ways that reduce her tax liabilities by choosing when to realize gains or losses. In this section we describe some models for optimal tax trading.

One of the earliest and most basic models for optimal tax trading was introduced by Constantinides (1983). In this model it is assumed that the tax rate on capital gains is independent of the length of the holding period. It is also assumed that capital losses generate tax rebates. Finally, it is assumed that there are no transaction costs, no capital loss restrictions, and no wash-sale restrictions. A wash sale occurs when an asset is sold at a capital loss and the same or substantially identical one is also purchased within 30 days before or after the sale. Under these assumptions the optimal tax-trading strategy is relatively simple: Realize losses as soon as they occur and defer gains indefinitely. By realizing losses, the investor gets a tax rebate. If the investor did not realize the loss as soon as it happened, the opportunity for a tax rebate could disappear. Constantinides's model can be extended to account for proportional transaction costs. If there are proportional transaction costs, then the optimal tax-trading strategy would still be to defer gains but to realize losses only beyond a certain threshold. The exact size of the threshold depends on the size of the transaction.

In a more elaborate follow-up article Constantinides (1984) proposed a model that considers a more realistic setting where the tax rate depends on the length of the holding period. In this model the sale of assets with long-term status is taxed at a rate lower than that of assets with short-term status. In this case the optimal tax-trading strategy still calls for realizing losses as soon as they occur. In addition and somewhat surprisingly, it is also sometimes optimal to sell (and immediately repurchase) assets with an embedded long-term gain. The rationale for this action is that there is a "re-start" option associated with resetting the tax basis and having the opportunity to realize short-term losses. The value of this re-start option depends on the asset volatility and the ratio of the short-term and long-term capital tax rates. The following example provided by C. Spatt [1] illustrates this phenomenon.

Example 14.1

Consider an asset with current price $P_0 = \$20$. Suppose that at dates $t = 0, 1$ we have

$$P_{t+1} = \begin{cases} P_t + k & \text{with probability } 0.5 \\ P_t - k & \text{with probability } 0.5. \end{cases}$$

Assume an investor buys one share of this asset at date $t = 0$. Our goal is to determine the trading strategy (realize/not realize) at dates $t = 1, 2$ that minimizes expected taxes.

Example 14.1

(a) First consider the following case. The short-term and long-term capital gain tax rates are respectively $\tau_s = 0.5$, and $\tau_\ell = 0.5y$, where $0 < y < 1$. The sale of shares held for one period can be treated as either short-term or long-term depending on what is more advantageous to the investor. The sale of shares held for two periods is treated as long-term. Assume there are no transaction costs.

In this case at dates $t = 1$ and $t = 2$ it is optimal to realize losses. At date $t = 1$ it is optimal to realize a long-term gain if $y < 0.5$. See Exercise 14.2.

(b) Now consider the case when the capital gain tax rate is $\tau = 0.2$ for both long-term and short-term gains or losses. Assume a transaction cost of 0.5 per share traded.

In this case at date $t = 2$ it is optimal to realize a loss if $k > 1.25$. At date $t = 1$ it is optimal to realize a loss if $k > 5$. Since short-term and long-term rates are the same, it is optimal not to realize gains at any date.

Portfolio Choice with Taxes

We now turn our attention to the problem of dynamic portfolio choice in the presence of capital gains taxes. The model below is a simplified version of a model proposed by Dammon et al. (2001).

We consider an economy with a risky and a risk-free asset where investors live for T periods. We also assume that in this economy investors are endowed with some initial capital and their goal is to maximize some expected utility of consumption C_t at dates $t = 0, 1, \ldots, T$ and bequest W_T at date T. The return of the risk-free asset between date $t-1$ and date t is r. The price of the risky asset is serially independent and follows a binomial process. Let P_t denote the price of the risky asset at date t. Let n_t and m_t denote respectively the number of shares of the risky and risk-free assets held right after trading at date t. Throughout the model we will assume no shorting, i.e., we will impose the constraints $n_t \geq 0$ and $m_t \geq 0$.

We assume that capital gains are taxed at a rate τ, and capital losses are credited at the same rate. To compute the capital gain triggered by an asset sale, we assume that the tax basis P_t^* for the shares at date t is the weighted average price of those shares. Therefore, the tax basis P_t^* evolves according to the following law of motion:

$$P_t^* = \begin{cases} \frac{n_{t-1} \cdot P_{t-1}^* + (n_t - n_{t-1})^+ \cdot P_t}{n_{t-1} + (n_t - n_{t-1})^+} & \text{if } P_{t-1}^* < P_t \\ P_t & \text{if } P_{t-1}^* \geq P_t \end{cases}$$

Right after trading at date t, the realized capital gain or loss G_t is given by

$$G_t = \begin{cases} (n_{t-1} - n_t)^+ \left(P_t - P_{t-1}^*\right) & \text{if } P_{t-1}^* \leq P_t \\ n_{t-1} \left(P_t - P_{t-1}^*\right) & \text{if } P_{t-1}^* \geq P_t. \end{cases}$$

We have the following inter-temporal balance of wealth equation that relates the portfolio holdings at dates $t-1$ to the portfolio holdings at dates $t = 1, \ldots, T-1$:

$$n_t P_t + m_t + C_t = n_{t-1} P_t + m_{t-1}(1+r) - \tau G_t$$

Similarly, at date T we have

$$W_T = n_{T-1} P_T + m_{T-1}(1+r) - \tau G_T$$

The portfolio choice problem can be stated as a dynamic programming problem where the state variables at date t are $\left(P_t, P_{t-1}^*, n_{t-1}, m_{t-1}\right)$, the actions at time t are (n_t, m_t, C_t) and the objective is

$$\max_{C_t, n_t, m_t} \mathbb{E}\left[\sum_{t=0}^{T} U\left(C_t, t\right) + B\left(W_T\right)\right].$$

The following example illustrates the striking effect of taxes in portfolio choice.

Example 14.2

Assume that at date $t = 0$ the holdings are $n_{-1} > 0$ and $m_{-1} = 0$. In other words, our entire portfolio is invested in the risky asset. Assume $r = 0, P_0 = 1, 0 < 1 - k < P^*_{-1} = 1 - \delta < 1$, and

$$P_1 = \begin{cases} P_0 + k \text{ with prob } 1/2 \\ P_0 - k \text{ with prob } 1/2 \end{cases}$$

Assume $T = 1$ and our goal is to determine the portfolio holdings at date $t = 0$ so as to maximize some utility of final wealth $\mathbb{E}\left(U\left(W_1\right)\right)$. Since there is no consumption at date $t = 0$ we have

$$n_0 P_0 + m_0 = n_{-1} P_0 + m_{-1} - \tau\left(n_{-1} - n_0\right)^+\left(P_0 - P^*_{-1}\right).$$

Thus

$$m_0 = \left(n_{-1} - n_0\right)\left(1 - \tau\delta\right)$$

Example 14.2

And so the only variable in our problem is n_0 subject to the constraints $0 \leq n_0 \leq n_1$.
Furthermore, the balance of wealth equation yields

$$\begin{aligned} W_1 &= n_0 P_1 + m_0 + \tau n_0 \left(P_0^* - P_1\right)^+ \\ &= \begin{cases} n_0(\tau\delta + k) + n_{-1}(1 - \tau\delta) & \text{with prob } 1/2 \\ n_0(\tau k - k) + n_{-1}(1 - \tau\delta) & \text{with prob } 1/2 \end{cases} \end{aligned}$$

It is evident that in the absence of taxes ($\tau = 0$), the optimal holding of the risky asset is $n_0 = 0$ for any positive level of risk aversion. However, it is easily checked numerically that n_0 may vary all the way between 0 and n_{-1} for positive values of τ. In particular, for $\tau = 0.2, \delta = 0.1, k = 0.2$ we get $n_0 = 0.8352$ for the logarithmic utility function.

The numerical solution to the more general model in Dammon et al. (2001) reveals the following interesting insights. As expected, it is optimal to realize capital losses as soon as they occur. Since diversification is more valuable to young investors, it is optimal for them to sell assets with large embedded capital gains to rebalance their portfolios. On the other hand, elderly investors defer most capital gains. Because in the US tax code there is a tax forgiveness at death, it is optimal for elderly investors to increase their allocations to equity as they approach their terminal age.

If in addition to some initial endowment an investor receives income, then the following insights are again revealed by a numerical solution to the model. Young investors hold more equity, very much in line with popular financial planning advice. Because of capital gain taxes, it is optimal to use income to adjust asset allocation instead of selling assets with embedded capital gains. In years immediately prior to retirement it is optimal to reduce equity allocation, again in line with popular financial planning advice. Finally, beyond retirement it is optimal to have a gradual increase in equity holdings.

Asset Allocation and Asset Location

The availability of various kinds of tax-deferred retirements accounts such as 401K, 403(b), IRA, and Keough give investors the ability to shelter some of their assets from taxes. Since assets may be held both in a tax-deferred as well as in a taxable account, the location decision has important implications on portfolio choice. Dammon et al. (2004) developed a model to study this problem. Via an arbitrage argument, they show that it is optimal to allocate assets to the tax-deferred account in descending order of tax exposure until the limit of the tax-deferred account is reached. In particular, the assets with the highest taxable yields, such as taxable bonds, should go in the tax-deferred account. If the limit of the tax-deferred account is reached, then assets with lower taxable yields should be allocated to the taxable account.

The numerical solution to the model in Dammon et al. (2004) also shows that, the larger the fraction of wealth in the tax-deferred account, the higher the fraction of total wealth allocated to assets with higher taxable yield.