Martingales and Portfolio decisions: a user's guide

Heinz Zimmermann

Abstract   This article gives an overview and introduction to the Martingale approach to multi-period (dynamic) portfolio decisions. While Martingale pricing techniques have long been used with considerable success in the pricing of derivatives and financial assets in general, their potential to improve the practice of dynamic portfolio decisions is not sufficiently recognized yet. This article shows that the approach is, in principle, not difficult to implement for readers equipped with standard option replication techniques if markets are sufficiently "complete" in order to provide investors with the relevant information about the pricing of financial risks. The article provides a practical guide to implement the basic features of the approach in a binomial framework.

H. Zimmermann
WWZ – Universität Basel
Abteilung Finanzmarkttheorie/Department of Finance
Holbeinstrasse 12, 4051 Basel, Switzerland
E-mail: heinz.zimmermann@unibas.ch
Tel.: +41-61-2673316
Fax: +41-61-2670898

1 Introduction

In the past decade, the theory of asset pricing made substantial progress both with respect to theoretical refinement and empirical testing. Compared to this, the advances in portfolio theory are much more limited. This is also reflected in the practice of asset management where single-period or myopic (i.e., short-sighted) portfolio decisions still dominate. At least when it comes to rebalancing decisions, the need for a more extensive theoretical framework is recognized. In this respect, the text of Campbell and Viceira (2002) by focusing on multi-period or live-cycle portfolio decisions attracted the interest of many investment professionals and researchers. Still, the theoretical and empirical advances of the asset pricing – call it "Martingale pricing" – literature are hardly exploited to improve dynamic portfolio decisions. Surprisingly, the respective theoretical foundations have been laid two decades ago, however, without much impact on the practice of investment management.

The papers by Cox and Huang (1989), Karatzas et al. (1987), Pliska (1986) and possibly others addressed the intertemporal portfolio selection problem in a completely new way: they recognized that in a complete market, where individuals have full information about the pricing of future states, the objective of the (possibly multi-period) investment decision can be formulated so as to determine the optimal distribution of final wealth. The portfolio selection problem can then be understood as a standard replication strategy familiar from option pricing.

The basics of the approach can be summarized by the insight that in complete markets, information can be exploited from financial markets so as to improve, or at least simplify, multi-period portfolio decisions – in the same way as the Martingale or arbitrage approach simplifies the pricing of assets. What is good for asset pricing models should not be worse for portfolio decisions!

While the original papers have developed the Martingale portfolio approach in a continuous-time stochastic framework with all analytical virtues, the practical implementation of the approach requires some additional thoughts. The textbook treatment of Campbell and Viceira (2002, chap. 5.2) also assumes a continuous time setting, while Cvitanic and Zapatero (2004) and Cerny (2004) contain discrete time (binomial) examples.

The paper is structured as follows: The basics of the methodology are illustrated in the first section using a simple numerical example and three equations only. In section 2, three Martingale results are presented which are key to understanding the approach. Sections 3 and 4 highlight how the Martingale approach is related to the classic dynamic portfolio opti-

mization problem. In section 5, a binomial implementation of the model is proposed. This is then illustrated by a worked-out numerical example in section 6. A short section concludes the paper. Although there are several formulae in the article, the reader should not be discouraged and notice that the emphasis is put on the implementation of the approach.

2 The Martingale Portfolio approach in a nutshell

In the classic approach to pricing derivatives, the terminal payoff of a security across all possible states s is given by its contractual features; we consider a call option on stock S with exercise price 90 maturing in two years. The stock price moves over 2 subsequent periods, starts at 100 and moves either up or down by 20% each period; so there are three states at the option's maturity (144, 96, and 64). Moreover, it is assumed that investors know the current price of one dollar accruing in each state $p_{0,T}^*(s)$ (called "state price").

The current, fair value of the contract can then be determined by multiplying payoffs with state prices:

$$C_0 = \sum_s X_T(s) \times p_{0,T}^*(s). \tag{1}$$

	Stock Price S_T	Payoff X_T	State Price $p_{0,T}^*$	$X_T \times p_{0,T}^*$
State 1 [uu]	144	54	0.35	19.13
State 2 [du]	96	6	0.43	2.55
State 3 [dd]	64	0	0.13	0.00
Sum				21.68

Why do investors know all state prices? They can be from observed prices of other financial instruments, or from a theoretical model. A market where state prices are uniquely determined is called "complete".

In a complete market, every possible distribution of cash-flows across future states can be manufactured (or replicated) with the existing financial instruments – either statically or by a dynamic strategy. In our example, we introduce a second asset, a risk-free one-period bond yielding 5%. Then the replicating strategy is as follows.

	$T = 0$	$T = 1$ Unbalanced	$T = 1$ Rebalanced		$T = 2$
					[uu]
				S	144.00
				B	−90.00
					54.00
			[u]		
S		92.14	120.00		
B		−57.86	−85.71		[ud]
		34.29	34.29	S	96.00
				B	−90.00
					6.00
Stocks	76.79				
Bonds	−55.10				
Total	21.68				[du]
			[d]	S	18.00
S		61.43	15.00	B	−12.00
B		−57.86	−11.43		6.00
		3.57	3.57		
					[dd]
				S	12.00
				B	−12.00
					0.00

Notice that this is a self-financing strategy – the rebalancing after the first period does not require any cash inflow or outflow.

Is this an optimal portfolio strategy? Would the investor get the maximum expected utility if she derived her entire future wealth from this specific option? The asset pricing literature typically does not address the question whether a specific future cash flow distribution (such as 54, 6, 0 in the example) is optimal for investors. The contractual characteristics of the specific security are regarded as given.

This contrasts the task of optimal portfolio selection, where we seek to manufacture an end-of-period distribution of wealth which is optimal in the sense of maximizing a specific objective function, such as expected utility.

In order to address the question whether the wealth distribution derived from the previous option contract is optimal, we must specify the utility function of the investor. If this is done, optimality requires that the marginal utility of wealth, weighted by the state probability $\pi_{0,T}(s)$ and normalized by the state price $p_{0,T}^*$, is the same across all states:

$$\frac{\pi_{0,T}(s_1) \times U'\left[W_T^{\text{opt}}(s_1)\right]}{p_{0,T}^*(s_1)} = \frac{\pi_{0,T}(s_2) \times U'\left[W_T^{\text{opt}}(s_2)\right]}{p_{0,T}^*(s_2)} = \cdots . \quad (2)$$

In a complete market, this condition uniquely determines the distribution of optimal wealth across all states. Assuming a utility function with constant relative risk aversion (CRRA γ, $U(W) = W^{1-\gamma}/(1-\gamma)$, $U'(W) = W^{-\gamma}$, and setting $\gamma = 5$, the figures in our numerical example become

	$X_T = W_T$	$\pi_{0,T}$	$\frac{\pi_{0,T} \times U'[W_T]}{p_{0,T}^*}$
State 1 [uu]	54	0.49	3.01E-09
State 2 [du]	6	0.42	1.27E-04
State 3 [dd]	0	0.09	∞

which demonstrates that the condition is largely violated; of course, no risk averse investor would accept a wealth of zero in any state. Therefore, the portfolio strategy is not optimal.

How can an optimal strategy be identified, given the level of initial wealth (21.68)? Of course, one has to start directly with the optimal wealth distribution in $T = 2$ and then derive the respective "replicating" strategy: the wealth distribution can be regarded as the payoff of a financial security and can be replicated correspondingly. This is the key idea of the Martingale portfolio approach.

The optimal wealth distribution can be directly derived from the optimality condition (2); the wealth ratio between any two states must satisfy

$$\frac{W_T(s_1)}{W_T(s_2)} = \left(\frac{p_{0,T}^*(s_1)/\pi_{0,T}(s_1)}{p_{0,T}^*(s_2)/\pi_{0,T}(s_2)}\right)^{-1/\gamma} \equiv \left(\frac{m_{0,T}(s_1)}{m_{0,T}(s_2)}\right)^{-1/\gamma} . \quad (3)$$

The state price-to-probability ratio is also called "stochastic discount rate" or "(state) deflator" and is denoted by $m_{0,T}(s)$. Optimality thus implies that the distribution of wealth is perfectly negatively correlated with the state deflator. If the deflator and the utility function of the investor are known, then the optimal distribution of wealth can be derived. Using wealth in state 2 to normalize, this becomes

| | State price $p_{0,T}^{*}$ | Probability $\pi_{0,T}$ | Deflator $m_{0,T}$ | Optimal wealth ratios | | Opt. Wealth W_T^{opt} |
				State 2	State 3	
State 1 [uu]	0.35	0.49	0.72	1.0696	1.1441	1.0696
State 2 [du]	0.43	0.42	1.01		1.0696	1.0000
State 3 [dd]	0.13	0.09	1.42			0.9349

Remark Most figures in this and in the following tables are displayed up to two decimal places only.

In order to generate the wealth distribution displayed in the last column, a current wealth of

$$W_0 = \sum_s W_T(s) \times p_{0,T}^{*}(s) = 0.9234$$

is required [by equation (1)]. Adjusting to the initial wealth of 21.68, the following optimal distribution results:

	Optimal wealth W_T^{opt}	State price $p_{0,T}^{*}$		Optimal wealth rescaled W_T^{opt}	
State 1 [uu]	1.0696	0.35	0.3790	25.1173	8.8993
State 2 [du]	1.0000	0.43	0.4252	23.4826	9.9841
State 3 [dd]	0.9349	0.13	0.1192	21.9544	2.8003
Current wealth			0.9234		21.6837

The final task is to find the replicating strategy which generates the optimal distribution:

	T = 0	T = 1 Unbalanced	T = 1 Rebalanced		T = 2
					[uu]
				S	4.90
				B	20.21
			[u]		25.12
		S 4.56	4.09		
		B 18.76	19.25		[ud]
		23.34	23.34	S	3.27
				B	20.21
Stocks	3.80				23.48
Bonds	17.89				
	21.68				[du]
			[d]	S	4.58
		S 3.04	3.82	B	18.90
		B 18.78	18.00		23.48
		21.82	21.82		
					[dd]
				S	3.06
				B	18.90
					21.95

It can be shown, that this is indeed a strategy satisfying the principle of dynamic optimality (formulated by Bellman and others). All that is needed is

1. a known distribution of state prices, or deflators;
2. a utility function based on which the optimal wealth distribution can be derived;
3. a recursive algorithm (equation 1) relating the terminal wealth to the current wealth (subsequently called Martingale wealth-constraint).

Then, the simple recursive replication algorithm known from option pricing can be applied to derive the optimal portfolio strategy.

Of course, the presented framework is too simplified to afford interesting applications and deeper insights. Therefore, some more details are presented in the next sections. However, the essential steps remain the same even if the stochastic and economic setting becomes more complex.

3 Martingale basics

3.1 State prices and deflators

In a complete market, the Arrow-Debreu or state prices in $t = 0$ for claims in $t = T, p_{0,T}^*$, can be uniquely determined from the observed prices of existing financial assets. In order to prevent arbitrage, the future state-dependent cash-flows of any security $X_{iT}(s)$ are related to the current asset value X_{i0} by

$$X_{i0} = \sum_s p_{0,T}^*(s)X_{iT}(s) = \sum_s \pi_{0,T}(s)m_{0,T}(s)X_{iT}(s) = E_0(m_T X_{iT}).$$

(4)

The definition of the deflator $m_{0,T}(s)$ was already shown in equation (3). The deflator makes it possible to write the valuation equation as (conditional) expectation. $\pi_{0,T}$ and $E_0[]$ are, respectively, the conditional probability and expectation based on the information available up to time $t = 0$. The discount factor $m_{0,T}$ must be the same for all assets in order to prevent arbitrage. Thus, the deflator relates financial assets' prices and expected future payoffs; it is therefore called the "pricing kernel" of the asset universe.

Alternative representations of equation (4) are also used. For instance, we can exploit the fact that the sum of all state prices equals the riskless discount factor, $\sum_s p_{0,T}^*(s) = \frac{1}{B_T}$, and write

$$X_{i0} = \sum_s p_{0,T}^*(s)B_T \frac{X_{iT}(s)}{B_T} \equiv \sum_s \hat{\pi}_{0,T}(s)\overline{X}_{iT}(s) \equiv \hat{E}_0(\overline{X}_{iT}), \quad (5)$$

where the weights $\hat{\pi}_{0,T}(s)$ add up to unity across all states, and can thus be understood as "new" probabilities. $\hat{E}[]$ denotes the mathematical expectation under the new probability measure; the bar symbol on a variable, e.g. $\overline{X}$, denotes riskless discounting,

$$\overline{X}_{it} \equiv X_{it}e^{-rt} \equiv \frac{X_{it}}{B_t}, \text{ hence } X_{i0} = \overline{X}_{i0}.$$

3.2 Discounted price processes as Martingales

In the following sections, the "Martingale property" of securities prices and wealth plays a key role. Under the original probabilities π one has

$$m_{0,t}X_{it} = E_t(m_{0,T} X_{iT}),$$

(6)

where $m_{0,0} = 1$. Hence, the discounted price process $\{m_t X_{it}\}$ is a Martingale using π. Under the new probabilities $\hat{\pi}$, we have

$$\overline{X}_{it} = \hat{E}_t (\overline{X}_{iT}). \tag{7}$$

Hence, the riskless discounted price process $\{\overline{X}_{it}\}$ is also a Martingale, however based upon $\hat{\pi}$. These are called "Martingale" probabilities. In the language of measure theory, the step from (6) to (7) is called a "transformation of probabilities" and is typically denoted by a specific operator y which can be characterized in the following way: We combine equations (4) and (5) to write

$$X_{i0} = \sum_s \pi_{0,T}(s) y_{0,T}(s) \frac{X_{iT}(s)}{B_T},$$

which requires

$$y_{0,T}(s) \equiv m_{0,T}(s) B_T = \frac{\hat{\pi}_{0,T}(s)}{\pi_{0,T}(s)}. \tag{8}$$

$\{y_{0,t}\}$ is called the "Radon-Nikoym process", and in the special case of continuous probability densities, the "Radon-Nikodym derivative" is dQ/dP, where Q and P represent probability measures corresponding to the original and transformed probabilities, respectively. This measure-theoretic notation is also used in the representation of conditional expectations, i.e.,

$$X_{i0} = E_0^P (m_{0,T} X_{iT}) = E_0^P (y_{0,T} \overline{X}_{iT}) = E_0^Q (\overline{X}_{iT}). \tag{9}$$

3.3 The Radon-Nikodym process as a Martingale

It follows from (8) that $\sum_s \pi_{0,T}(s) y_{0,T}(s) \equiv \sum_s \hat{\pi}_{0,T}(s) = 1$, and since $y_{0,t} = 1$ given the information in t, we get

$$y_{0,t} = E_t (y_{0,T}), \tag{10}$$

which means that the Radon-Nikodym process is a Martingale as well.

3.4 Self-financing Portfolio strategies as Martingales

The Martingale properties of

  – discounted price processes [equations (6) and (7)] and
  – the Radon-Nikodym process [equation (10)]

imply a third result, namely, that (discounted) wealth is a Martingale under a self-financing strategy as well. Wealth is subsequently defined as the value of a portfolio consisting of two assets: a risky asset (stock) with price X_T, and a riskless bond. If δ_{T-1} denotes the number of stocks in the portfolio at the beginning of time interval $(T-1, T)$, the self-financing condition imposes the following restriction on the dynamics of wealth:

$$W_T = \delta_{T-1}X_T + (W_{T-1} - \delta_{T-1}X_{T-1})B_1. \tag{11}$$

It is shown in the Appendix and in the quoted literature, that the two Martingale properties plus (11) imply the Martingale property of discounted wealth:

$$m_{0,t}W_t = E_t\big(m_{0,T}W_T\big) \text{ resp. } y_{0,t}\overline{W}_t = E_t\big(y_{0,T}\overline{W}_T\big). \tag{12}$$

This will be called "Martingale wealth condition", and it will be shown in section 4 that it permits an alternative and much more elegant representation of multi-period portfolio selection problems than standard approaches.

4 Deflators and the optimal wealth distribution

The key insight of the Martingale portfolio approach is that in complete markets, individuals know everything about the stochastic deflator (or equivalently, about the state prices or the Radon-Nikodym process) – or are at least able to make reasonable estimates. This fact can be exploited to infer individuals' optimal wealth distribution at any future date. The reason is that the optimal portfolio decision of the individuals – and thus their optimally accumulated future wealth – is directly affected by the deflators [see equation (3)]. Therefore, the distribution of the deflator across states is directly related to the risk-preferences, i.e., the shape of the utility function, of the individuals. Of course, the exact relationship depends on the specific optimization framework being used. In a complete market, however, there is not much choice in this respect because any portfolio selection problem can be reduced to a portfolio allocation problem across state securities.

Let us denote the number of state securities held in the portfolio by $\underline{\delta}^*$. Maximizing the expected utility of end-of-period wealth under a budget constraint determined by the existing state prices leads to the Lagrangian

$$\sup_{\underline{\delta}^*,\theta} \sum_s \pi_{0,T}(s)\, U\left[W_T(s) = \delta^*(s)\right] - \theta \left\{ \sum_s \delta^*(s)\, p_{0,T}^*(s) - W_0 \right\} \tag{13}$$

with the first order condition

$$\frac{\pi_{0,T}(s)\, U'[W_T(s)]}{p_{0,T}^*(s)} = \frac{U'[W_T(s)]}{m_{0,T}(s)} = \theta \tag{14a}$$

to hold for each state; this is the optimality condition used in the numerical example of section 1 [equation (2)]. Taking expectations yields

$$\frac{\sum\limits_s \pi_{0,T}(s)\, U'[W_T(s)]}{\sum\limits_s \pi_{0,T} m_{0,T}(s)} = \frac{E_0^P\left[U'(W_T)\right]}{E_0^P\left[m_{0,T}\right]} = B_T\, E_0^P\left[U'(W_T)\right] = \theta, \tag{14b}$$

and combining (14a) and (14b) gives

$$m_{0,T} = \frac{1}{B_T} \frac{U'\left(W_T^{\text{opt}}\right)}{E_0^P\left[U'\left(W_T^{\text{opt}}\right)\right]}, \tag{14c}$$

which shows the relation between the deflator, the shape of the (marginal) utility function, and optimal wealth of the individual. Of course, an explicit solution for the optimal wealth distribution across states requires a specification of the utility function.

5 The Martingale Portfolio approach

Consider the simplest case of the multi-period portfolio selection problem, where we accumulate our wealth over a finite number of periods with intermediate rebalancing, but without withdrawals, investments or consumption. The objective is to maximize the expected utility of the accumulated wealth by selecting a sequence of portfolio strategies over time, denoted by a vector

$$\underline{\delta}^{t,T-1} \equiv \left[\delta_t, \delta_{t+1}, \ldots, \delta_{T-1}\right],$$

where the individual elements δ_j represent the portfolio holdings (in the simplest case, the number of risky assets in a two-asset portfolio with

riskless borrowing and lending) in the time interval $(j, j + 1)$. The optimization problem is thus

$$\sup_{\underline{\delta}^{0,T-1}} E_0^P \left[U\left(\tilde{W}_T \right) \right]$$

plus a standard budget restriction on the current wealth W_0.

5.1 The Bellman principle of dynamic optimality

The classical solution to this problem is the Bellman principle of dynamic programming, which assumes that the investor knows the optimal portfolio strategy one period ahead, $\underline{\delta}^{1,T-1}$, and is thus able to compute the maximum expected utility he can derive from his current wealth:

$$J_1 = \sup_{\underline{\delta}^{1,T-1}} E_1^P \left[U\left(\tilde{W}_T^{\text{opt}} \right) \right].$$

J_t is called "indirect utility function". The principle of dynamic optimality states that the current portfolio decision can be characterized by

$$J_0 \equiv \sup_{\delta_0} E_0^P \left[J_1, W_1^{\text{opt}} \right] \tag{15}$$

which means that the optimal strategy one period ahead directly affects the current (optimal) decision. This makes it possible to replace the full sequence of portfolio decisions, as characterized by the vector $\underline{\delta}^0$, by a single portfolio decision, δ_0.

Of course, the optimal strategy one period ahead is not known either – but it can be derived from the optimal decision two periods ahead in exactly the same way. This can be repeated until we reach the final time interval $(T - 1, T)$. However, in T, no optimization is required any more. Technically speaking, the law of iterated expectations makes it possible to substitute conditional expected values according to

$$E_0^P [W_T] = E_0^P \left[E_{T-1}^P (W_T) \right],$$

or

$$E_0^P [W_T] = E_0^P \left[E_{T-2}^P (W_T) \right] = E_0^P \left[E_{T-2}^P \left[E_{T-1}^P (W_T) \right] \right],$$

and so forth until $E_0^P [W_T] = E_0^P \left[E_1^P (W_T) \right]$. This procedure also applies to the optimally invested wealth. As a consequence, the Bellman principle stipulates that the optimal portfolio strategy can be found by a recursive (or backward) algorithm.

In summary, the Bellman principle makes it possible to substitute a full sequence of portfolio decisions by a sequence of one-period decisions.

However, there is an even more economical way to solve the portfolio problem — based on the Martingale property of the discounted wealth process (12) and the assumption of complete markets.

5.2 The Martingale solution

It can be shown that the multi-period optimization problem can be represented by a one-period problem: in this setting, the investor directly determines the optimal distribution of his final (end-of-period) wealth W_T^{opt} while facing the Martingale wealth condition as a (quasi static) budget constraint. We thus have the optimization problem

$$\sup_{\tilde{W}_T, \theta} E_0^P \left[U \left(\tilde{W}_T \right) \right] - \theta \left\{ E_0^P \left[m_{0,T} \tilde{W}_T \right] - W_0 \right\}, \tag{16}$$

where θ is the Lagrange multiplier. Apparently, the problem looks like a one-period optimization problem!

A standard dynamic replication strategy (known, for example, from option pricing theory) can then be determined to replicate the optimally distributed final wealth of the investor. The replication strategy is identical to the optimal dynamic portfolio policy. The technical equivalence with the original portfolio problem relies on convex duality theory in optimization; these methods were introduced in a continuous-time setting by Karatzas et al. (1987) and Cox and Huang (1989). Recent textbook treatments can be found in, for example, Cerny (2004, chap. 9.4) or Cvitanic and Zapatero (2004, chap. 4.4).

The first order conditions for the problem stated in equation (16) are

$$U' \left(W_T^{\text{opt}}(s) \right) = \theta \, m_{0,T}(s) \quad \forall s, \tag{17a}$$

and

$$E_0^P \left[U' \left(W_T^{\text{opt}} \right) - \theta m_{0,T} \right] = 0. \tag{17b}$$

These conditions are identical to those derived in the one-period complete markets setting (14a) and (14b).

Obviously, solving for the optimal wealth across states, $W_T^{\text{opt}}(s)$, requires inversion of the marginal utility function. For computational purposes, this requires a specification of the utility function. We will again use CRRA in the following section.

Moreover, the Lagrange multiplier θ must be explicitly determined. It can be directly derived from the Martingale wealth condition $W_0 = E_0^P \left[m_{0,T} \tilde{W}_T \right]$, if $\tilde{W}_T$ is substituted by the optimal solution from (17a), which is a pure function of $m_{0,T}$ and θ:

$$W_0 = E_0^P \left[m_T \, f \left(\theta, m_{0,T} \right) \right].$$

This is an equation with one unknown, θ, because the $m_{0,T}$ are completely specified in a complete market. Substituting the solution for θ, θ^{opt}, back in the system of equations (17a), we can now explicitly solve for $W_T^{\text{opt}}(s)$ across all states. The example in the next section will highlight this procedure.

The economic interpretation of the multiplier can be directly derived from (17b),

$$\theta = \frac{E_0^P\left[U'\left(W_T^{\text{opt}}\right)\right]}{E_0^P\left(m_{0,T}\right)} = B_T\ E_0^P\left[U'\left(W_T^{\text{opt}}\right)\right],$$

which is the expected marginal utility of optimally invested wealth compounded at the riskless rate; see also (14b) where the same interpretation holds in the state-preference setting.

The Martingale approach represents a drastic simplification of the multi-period portfolio problem – but it comes at a cost: the assumption of complete markets. However, this assumption has proved its merits in asset pricing; so why not exploit it for designing dynamic portfolio strategies? In reality markets are neither complete nor incomplete. The degree of (in)completeness relies on how well arbitrary claims can be replicated with existing securities.

In the following, we want to show how this approach can be implemented in practice based on simple binomial dynamics of asset prices and wealth. In contrast, many of the technical proofs of the original papers apply to continuous time economies. While continuous-time models are analytically more appealing, the binomial setting offers a simple way in implementing the approach.

6 A specific analytical example

The general characterization of the Martingale approach in the previous section gets a much more intuitive appeal if several specific assumptions are made:

(a) The portfolio: we assume that the investor faces a single risky asset (the "market" portfolio) and a one-period riskless asset. In principle the riskless rate may change over time; but we will assume that it is constant.

(b) We assume that the risky asset value in $(t+1)$ and hence the accumulated wealth, is characterized by a binomial process with up-probability $\pi_{t,t+1}$ and down-probability $1-\pi_{t,t+1}$, conditional on the information in t. The probabilities can, in principle, change over time.

(c) The utility function: we assume a function implying CRRA,

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}$$

where γ is the coefficient of relative risk aversion. Marginal utility is then $U'(W) = W^{-\gamma}$. This is an assumption which applies to a specific investor, not to the market as a whole, so that it does not impose a restriction on the evolution of the stochastic process of the deflator or on the market risk premia of the risky asset.

(d) Market completeness: a complete market can be characterized in several ways, for example by the Radon-Nikoym process $\{y_{0,t}\}$, the deflator process $\{m_{0,t}\}$, or the process of state prices $\{p_{0,t}^*\}$ depending on which specification is the most convenient. This may, however, be too abstract for practical purposes. Instead, we assume that the investor is able to determine the process of the so called "market price of risk", denoted by $\{\lambda_{0,t}\}$ for the time interval $(t-1, t)$; it may apparently change over time, but it is "measurable" with respect to the information available in t.

Analytically, the market price of risk is defined as the standard deviation of the deflator. In more practical terms, it is the "Sharpe Ratio" of the market portfolio, which is by definition perfectly negatively correlated with the deflator. This approach seems to be justified because most investment practitioners are familiar with the Sharpe Ratio of their portfolios and the associated benchmarks; moreover, many quantitative tactical asset allocation tools directly focus the determinants of the market price of risk.

The Martingale portfolio selection procedure can be broken up into three consecutive steps:

1. modeling the Radon-Nikodym process, or the deflator process;
2. deriving the terminal distribution of optimal wealth;
3. determining the optimal replication/portfolio strategy.

There is also an intermediate step 1^{bis}, the modeling of the stochastic process of the underlying (risky or market) security.

6.1 Step 1 the stochastic deflator process

Given the binomial probabilities $\pi_{0,1}$ and $1 - \pi_{0,1}$ and the market price of risk $\lambda_{0,1}$ for the first holding period, it can be shown that the adjusted or Martingale probabilities can be simply computed by

$$\hat{\pi}_{0,1}(s) = \pi_{0,1}(s) - \lambda_{0,1} \sqrt{\pi_{0,1}(s)\left(1 - \pi_{0,1}(s)\right)} \tag{18}$$

for the two binomial end-of-period states; see Zimmermann (2005, chap. 5). Computing forward, the entire process $\{\hat{\pi}_{t,t+1}\}$ can be derived. Based on this, we compute the probabilities of the entire tree conditional on the information in $t = 0$, $\{\hat{\pi}_{0,t}\}$. This is the basis for the Radon-Nikodym-process and the stochastic deflator process:

$$y_{0,t}(s) = \frac{\hat{\pi}_{0,t}(s)}{\pi_{0,t}(s)},$$

$$m_{0,t}(s) = \frac{y_{0,t}(s)}{B_t}.$$

6.2 Step 1$^{\text{bis}}$ the securities process

In order to derive the portfolio strategy in step 3, one also needs the specification of the stock price process of the risky asset $\{S_t\}$ which is consistent with the preceding assumptions. We proceed in the standard way known from binomial option pricing; we assume that the stock price follows a multiplicative binomial random walk

$$S_{t+1}(s^+) = S_t \times u_{t,t+1}, \; S_{t+1}(s^-) = S_t \times d_{t,t+1} \tag{19}$$

where $u_{t+1} (d_{t+1})$ is the realized "one plus return" in the upstate s^+ (downstate s^-) over time interval $(t, t + 1)$. We further assume, in the Cox-Ross-Rubinstein manner, that the returns are inversely related by

$$u = \frac{1}{d} = e^{\sigma}$$

where σ is approximatively the one-period standard deviation based upon π. In $t = 0$, the well-known no-arbitrage condition which determines the Martingale up-probability is

$$\hat{\pi}_{0,1}(s) = \frac{B_1 - d_{0,1}}{u_{0,1} - d_{0,1}},$$

which can then be used to solve

$$d_{0,1} = \frac{+B_1 \pm \sqrt{B_1^2 - 4(1 - \hat{\pi}_{0,1})\hat{\pi}_{0,1}}}{2(1 - \hat{\pi}_{0,1})} \tag{20}$$

where typically the negative sign in the numerator gives the appropriate solution. Proceeding that way forward, the entire process $\{S_t\}$ can be derived. Notice that, except for special cases, the derived binomial tree is non-recombining, which means that the computational task can be substantial if the number of time steps is high.

6.3 Step 2 the optimal wealth process

Based on the stochastic deflator process derived in Step 1 and the assumed CRRA utility function, we can derive the optimal wealth distribution across all states $W_T^{\text{opt}}(s)$. We take the system of equations (17a) to write

$$\left(W_T^{\text{opt}}\right)^{-\gamma}(s) = \theta \, m_{0,T}(s) \quad \forall s$$

with the solution

$$W_T^{\text{opt}}(s) = \theta^{-1/\gamma} \, m_{0,T}^{-1/\gamma}(s) \tag{21a}$$

The Lagrange multiplier θ still must be determined. We replace (21a) in the Martingale-wealth condition (12) to get

$$W_0 = E_0^P\left[m_{0,T} \, \theta^{-1/\gamma} \, m_{0,T}^{-1/\gamma} \right]$$

$$= E_0^P\left[m_{0,T}^{1-1/\gamma} \, \theta^{-1/\gamma} \right] = \theta^{-1/\gamma} E_0^P\left[m_{0,T}^{1-1/\gamma} \right],$$

which can be solved for the multiplier

$$\theta^{-1/\gamma} = \frac{W_0}{E_0^P\left[m_{0,T}^{1-1/\gamma} \right]}$$

and re-inserted in (21a) to obtain

$$W_T^{\text{opt}}(s) = W_0 \frac{m_{0,T}^{-1/\gamma}(s)}{E_0^P\left[m_{0,T}^{1-1/\gamma} \right]}, \tag{21b}$$

where the ratio can be interpreted as "one plus the rate of return" on the optimally invested wealth; a similar expression can be found in Cerny (2004, p. 210). It is entirely determined by the deflator and the risk aversion coefficient! This is the key result of the Martingale portfolio approach. All the remaining tasks can be fully adapted from the binomial option pricing approach.

The distribution of optimal terminal wealth $W_T^{\text{opt}}(s)$ can be regarded as the final payoff of a derivative security, so that the entire optimal wealth process before the terminal date, $\left\{W_t^{\text{opt}}\right\}$, $0 \le t < T$, can be recursively determined by using the Martingale wealth property $\overline{W}_{i,t-1} =$

$E^Q_{t-1}\left(\overline{W}^{\text{opt}}_{i,t}\right):$

$$W^{\text{opt}}_{t-1} = \frac{B_{t-1}}{B_t} E^Q_{t-1}\left(W^{\text{opt}}_t\right)$$

$$= \frac{\hat{\pi}_{t-1,t}\left(s^+\right) W^{\text{opt}}_t\left(s^+\right) + \left[1 - \hat{\pi}_{t-1,t}\left(s^-\right)\right] W^{\text{opt}}_t\left(s^-\right)}{1 + R}.$$

This corresponds to the recursive risk-neutral pricing procedure for derivatives. Unlike this case, however, the original wealth W_0 is given here.

6.4 Step 3 the replication strategy

The natural last step is to determine the multi-period portfolio strategy which perfectly replicates the optimal wealth distribution $W^{\text{opt}}_T(s)$ and the associated process $\left\{W^{\text{opt}}_t\right\}$, $0 \le t \le T$. This will be the optimal strategy. As for the derivatives in a complete market, the replication strategy can be directly derived from the binomial model; in t, the replicating portfolio includes in each state the following number of shares

$$\delta_t = \frac{W^{\text{opt}}_{t+1}\left(s^+\right) - W^{\text{opt}}_{t+1}\left(s^-\right)}{S_{t+1}\left(s^+\right) - S_{t+1}\left(s^-\right)},$$

and the amount invested in the riskless security is

$$\frac{u_{t,t+1}\, W^{\text{opt}}_{t+1}\left(s^-\right) - d_{t,t+1}\, W^{\text{opt}}_{t+1}\left(s^+\right)}{\left(u_{t,t+1} - d_{t,t+1}\right) B_1}.$$

Of course, the replication can be extended to more complex portfolios, including dividend payments (or consumption, liabilities) or transaction costs.

7 Numerical implementation

In this section, we illustrate the procedure outlined before by a concrete numerical example. Notice that CRRA was used for illustrative purposes and does not produce the most exciting results in studying multi-period portfolio decisions. If probabilities and the market price of risk are constant over time, CRRA does not induce investors to optimally adjust their portfolio weights (here: the stock–bond-mix) if wealth changes. Also, if the average investor had CRRA, the market risk premium would not fluctuate either. Therefore, the way to read the following examples is that we are studying the optimal asset allocation of an *individual* investor

Table 1 Constant market price of risk: optimal terminal wealth distribution and asset allocation distribution

Final states	Coefficient of relative risk aversion (γ)					
	1	2	3	4	5	10
uuu	368.36	215.87	175.99	158.12	148.06	129.35
uud	166.49	145.13	135.06	129.65	126.31	119.48
udu	166.49	145.13	135.06	129.65	126.31	119.48
udd	75.25	97.57	103.65	106.31	107.76	110.35
duu	166.49	145.13	135.06	129.65	126.31	119.48
dud	75.25	97.57	103.65	106.31	107.76	110.35
ddu	75.25	97.57	103.65	106.31	107.76	110.35
ddd	34.01	65.60	79.54	87.16	91.94	101.93
$E_0^P\left(W_T^{\mathrm{opt}}\right)$	175.34	141.62	131.30	126.36	123.47	117.86
Stocks	74.96%	37.51%	24.87%	18.58%	14.82%	7.37%
Bonds	25.04%	62.49%	75.13%	81.42%	85.18%	92.63%

Constant relative risk aversion (CRRA); starting wealth 100; market price of risk (Sharpe Ratio of market portfolio) 0.4; interest rate 4%; 3 periods

with CRRA in a *market* framework where the market risk premium is fluctuating.

Throughout the examples, we also assume a discrete interest rate of 4%; a constant value is, of course, not necessarily realistic given a time-varying market price of risk. We moreover assume that the original up- and down-probabilities are constant, 0.6 and 0.4.

We first assume a constant ex ante market price of risk, $\lambda = 0.4$, and a starting wealth of 100; the resulting optimal distribution of terminal wealth is displayed in Table 1 for alternative levels of relative risk aversion. Individuals choose a constant relative allocation of stocks and bonds over time, maintained by rebalancing after each period, as displayed at the bottom of the table.

As expected, by increasing the relative risk aversion, the distribution of optimal wealth narrows and the expected level of optimal wealth thereby decreases.

In the remaining part of this section, we present a worked-out numerical example.

Step 1 The stochastic deflator process

We assume a process for $\{\lambda_{o,t}\}$ with high values in bad states and low values in good states:

Market price of risk

Notice that these values are known at the beginning of each period; thus, 0.2 is the market price of risk in the third period after the market has moved up twice. A realistic process would perhaps not be so volatile over annual time intervals. The implied process of Martingale probabilities conditional on the information one period ahead, $\{\hat{\pi}_{t,t+1}\}$, is

---

Martingale probabilities
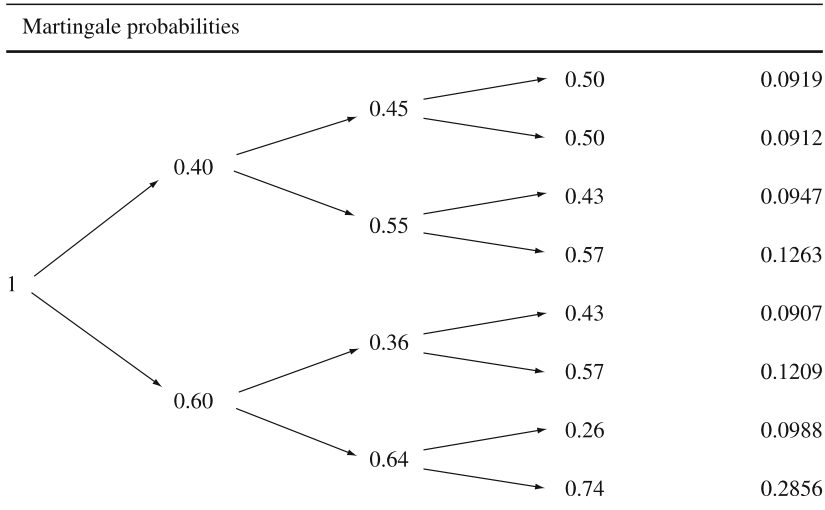
The terminal state probabilities conditional on the information in $t = 0$, $\hat{\pi}_{0,T}(s)$, are displayed in italics at the end of the tree. The implied Radon-Nikodym process $\{y_{0,t}\}$ is then

Radon-Nikodym

The derived stochastic deflator process $\left\{m_{0,t}\right\}$ is

State deflator

| Money mkt | 1.0400 | 1.0816 | 1.1249 |
| Riskless disc | 0.9615 | 0.9246 | 0.8890 |
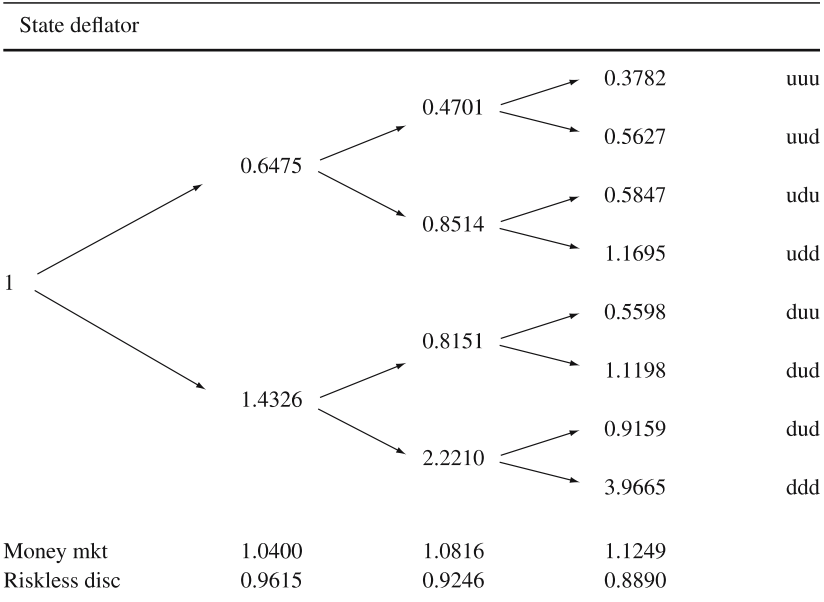
The money market process $\{B_t\}$ and its inverse, the riskless discount factors, are displayed at the bottom. We also present the time-path of each final state at the end of the tree. Unfortunately, this does not give an ideal (uniform) ranking of the states from "good" at the top (low deflator) to "bad" at the bottom (high deflator). Since the tree is non-recombining, it is not possible to ideally match both representations.

Step 1^{bis} The securities process

We assume a single risky asset and derive the up and down returns consistent with $\{\hat{\pi}_{t,t+1}\}$ and the interest rate according to equation (20); this gives

---

Up and down "one plus returns" (one step)

```
                                              1.3204
                            1.4797
                                              0.7574
            1.7129
                                              1.5861
                            0.6758
                                              0.6305
1
                                              1.5861
                            2.0377
                                              0.6305
            0.5838
                                              3.1190
                            0.4907
                                              0.3206
```

Notice that the process is not recombining. Assuming an initial stock price of 1 leads to the following price lattice $\{S_t\}$:

Stock price lattice

```
                                              3.3465
                            2.5346
                                              1.9196
            1.7129
                                              1.8361
                            1.1576
                                              0.7298
1.0000
                                              1.8869
                            1.1897
                                              0.7500
            0.5838
                                              0.8936
                            0.2865
                                              0.0919
```
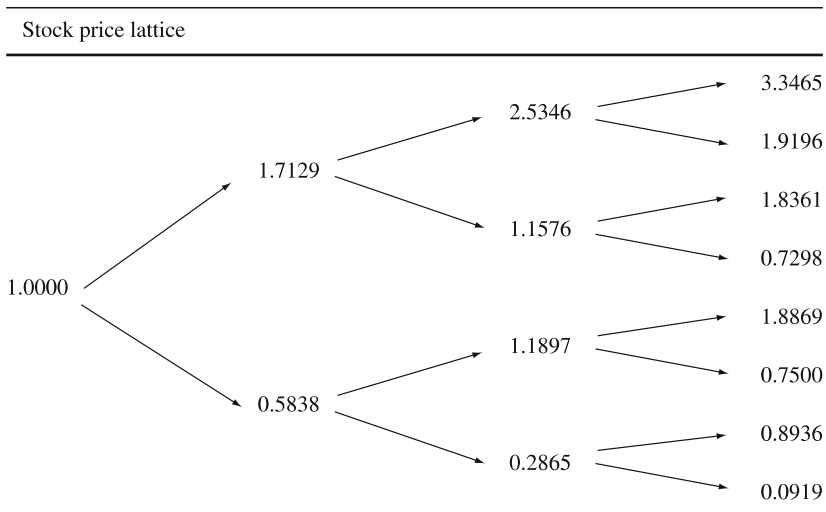

The expected one, two and three period returns (continuously compounded) are 10.76, 21.16 and 31.23%, respectively, which is only slightly less than time-proportional.

Step 2 The optimal wealth process
    Based on the terminal values of the state deflator, and assuming a relative risk aversion of 3, the expectation $E_0^P \left[m_T(s)^{1-1/\gamma} \right] = 0.8705$ can be computed and equation (21b) is used to derive the optimal wealth level in all terminal states; see the figures in italics:

Optimal wealth proces
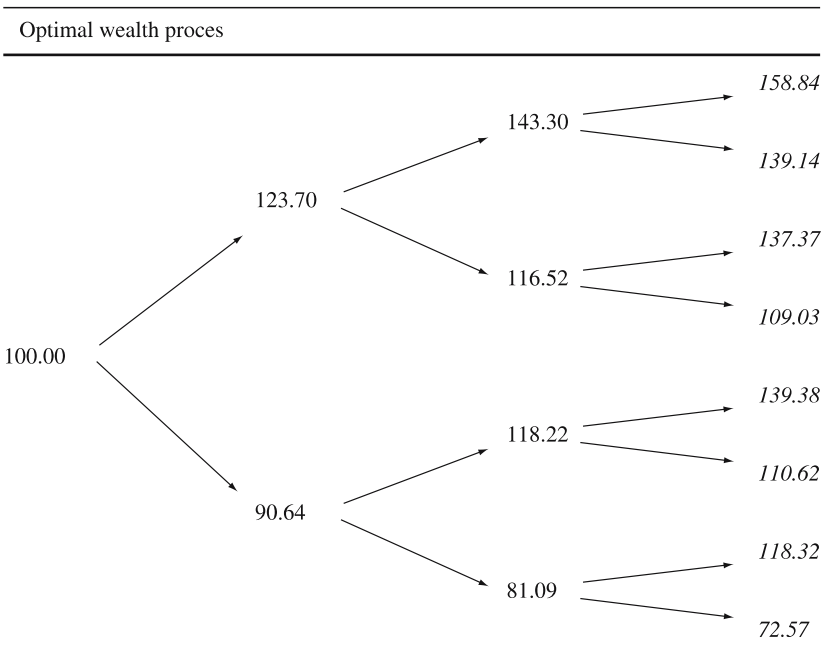
The entire optimal wealth process is then derived based on the Martingale wealth property represented by equation (22), which is the standard recursive risk-neutral pricing procedure. Obviously, the predetermined initial wealth of 100 results.

Step 3 The replication strategy
    The final step is to replicate the optimal wealth process. Again, the standard replication strategy from option pricing can be applied:

Absolute holdings

```
                                                                    35.00
                                                                   108.29
                                            33.32                  143.30
                                            90.38
                                           123.70                   29.66
                                                                    86.85
Stocks              29.28                                          116.52
Bonds               70.72
Total              100.00                                           30.09
                                                                    88.12
                                            24.00                  118.22
                                            66.64
                                            90.64                   16.35
                                                                    64.74
                                                                    81.09
```
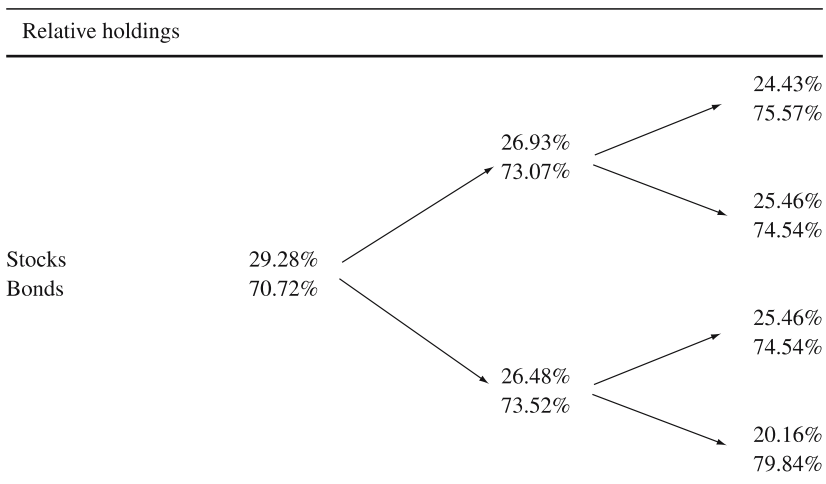

The lattice shows the dollar amounts optimally allocated to stocks and riskless bonds at the beginning of the first, second and third time period. It can easily be shown that the strategy is self-financing. It is interesting to notice that the absolute amount invested in stocks is rather stable — except in the first and second downstate. The relative amounts may be more interesting:

Relative holdings

```
                                                                  24.43%
                                                                  75.57%
                                          26.93%
                                          73.07%
                                                                  25.46%
                                                                  74.54%
Stocks              29.28%
Bonds               70.72%
                                                                  25.46%
                                                                  74.54%
                                          26.48%
                                          73.52%
                                                                  20.16%
                                                                  79.84%
```

The relative holding of stocks is decreased for a shrinking time horizon across all states, but the effect is moderate. Moreover, the decreasing market price of risk induces investors to decrease their stockholdings in the upmarket; for example, if the market price of risk were zero in the *uu*-state, then there would be no optimal holding of any stock in the last period.

Analyzing the portfolio effects in the downmarket is more tricky. If the market price of risk increases in bad states, then so does the deflator in these states (the bad state appears even worse), and thus the optimal level of wealth decreases. In terms of replication, this means that the "safety cushion" an investor faces at the beginning of the investment period appears more comfortable, which can be accomplished by taking some more risk. For example, if the market price of risk in the *dd*-state (0.7) were replaced by 1.0, the original stock holding would be 32% instead of 28%. But the adjustment of the replicating portfolio is difficult to characterize in general because the strategy naturally depends on how λ affects the pricing of the underlying risky asset in the various states. Moreover, the effects crucially depend on the degree of relative risk aversion.

Overall, the example highlights the potential of the Martingale methodology as a practical tool for analyzing the determinants of multi-period portfolio decisions in a binomial setting.

8 Conclusion

The article provides an overview and introduction to the Martingale portfolio approach and the way it can be implemented in a binomial framework. While Martingale pricing techniques have long been used with considerable success in the pricing of derivatives and financial assets in general, their potential to improve the practice of multi-period portfolio decisions is not sufficiently recognized yet.

Moreover, understanding multi-period portfolio decisions as optimal wealth-replication strategies also has the potential to fill the theoretical gap between portfolio management, financial planning, and the structuring of optimal financial products.

Of course, the practical applicability of the approach depends on the amount and the reliability of information that can be extracted from financial market prices to determine the pricing of the relevant financial risks investors are faced with. The degree of market completeness is a major challenge for the practical validity of the approach. Implementing the model in incomplete markets, taking into account transaction costs or other frictions, complicates the practical implementation considerably.

A discussion of these issues can, for example, be found in Eggers (2004). This article helps to understand the assumptions and basic steps in implementing the approach, as well as its potential and limitations. Whether it actually produces superior investment results compared to traditional strategies depends on the availability and quality of the relevant inputs.

Appendix

The self-financing condition of a portfolio with two assets, a risky stock with price X_T and a riskless bond with return r, $B_T = \mathrm{e}^{rT}$, can be characterized by

$$W_T = \delta_{T-1} X_T + (W_{T-1} - \delta_{T-1} X_{T-1}) B_1,$$

or in discounted terms

$$\overline{W}_T \equiv \frac{W_T}{B_T} = \delta_{T-1} \frac{X_T}{B_T} + \left(\frac{W_{T-1}}{B_T} - \delta_{T-1} \frac{X_{T-1}}{B_T} \right) B_1$$

$$= \delta_{T-1} \overline{X}_T + \left(\overline{W}_{T-1} - \delta_{T-1} \overline{X}_{T-1} \right).$$

We want to address the following expression:

$$E_t \left(y_{0,T} \overline{W}_T \right) = E_t \left[y_{0,T} \delta_{T-1} \overline{X}_T + y_{0,T} \overline{W}_{T-1} - y_{0,T} \delta_{T-1} \overline{X}_{T-1} \right];$$

applying the law of iterated expectations gives

$$E_t \left(y_{0,T} \overline{W}_T \right) = E_t \left[E_{T-1} \left(y_{0,T} \delta_{T-1} \overline{X}_T + y_{0,T} \overline{W}_{T-1} - y_{0,T} \overline{X}_{T-1} \right) \right]$$

$$= E_t \left[\delta_{T-1} \underbrace{E_{T-1} \left(y_{0,T} \overline{X}_T \right)}_{y_{0,T-1} \overline{X}_{T-1}} + \overline{W}_{T-1} \underbrace{E_{T-1} \left(y_{0,T} \right)}_{y_{0,T-1}} - \delta_{T-1} \overline{X}_{T-1} \underbrace{E_{T-1} \left(y_{0,T} \right)}_{y_{0,T-1}} \right]$$

$$= E_t \left[y_{0,T-1} \overline{W}_{T-1} \right]$$

and proceeding backwards until

$$E_t \left(y_{0,T} \overline{W}_T \right) = E_t \left[y_{0,t+1} \overline{W}_{t+1} \right] = y_{0,t} \overline{W}.$$

shows the Martingal property for $\left\{ y_{0,t} \overline{W}_t \right\}$ or $\left\{ m_{0,t} W_t \right\}$, respectively.

References

Campbell, J.Y., Viceira, L.: Strategic asset allocation. Portfolio choice for long term investors, Clarendon Lectures in Economics. New York: Oxford University Press, 2002

Cerny, A.: Mathematical techniques in finance. Princeton: Princeton University Press, 2004

Cox, J., Huang, C.-F.: Optimal consumption and portfolio policies when asset prices follow a diffusion process. J Econ Theory **49**, 33–83, (1989)

Cvitanic, J., Zapatero, F.: Introduction to the economics and mathematics of financial markets. Cambridge: M.I.T. Press 2004

Eggers, R.: Martingaleansätze in der Portfolioselektion. Dissertation Universität St. Gallen HSG, 2004

Karatzas, I., Lehoczky, J., Shreve, S.: Optimal portfolio and consumption decisions for a small investor on a finite horizon. SIAM J Control Optim **25**, 1157–1186 (1987)

Pliska, St.: A stochastic calculus model of continuous time trading: optimal portfolios. Math Oper Res **11** (1986)

Zimmermann, H.: Finance derivatives. Zürich: NZZ-Verlag 2005

About the Author

Heinz Zimmerman holds the Chair of Finance at the Faculty of Economics and Business of the University of Basel. He also heads the Department of Finance at the Wirtschaftswissenschaftliches Zentrum (WWZ). He is also dean of the faculty. Before, he was a professor at and a cofounder of the Swiss Institute of Banking and Fi-nance at the University of St. Gallen.

He has published about 180 scientific papers in books or journals, among these "Global Asset Allocation" (Wiley, 2003), and "Finance derivatives" (NZZ-Verlag, 2005). Heinz Zimmermann is a member of the editorial board of various international journals and is a member of the editorial board as well as a former editor-in-chief of Financial Markets and Portfolio Management.