Economic Notes by Banca Monte dei Paschi di Siena SpA,
vol. 33, no. 1-2004, pp. 55–81

Optimal Dynamic Trading Strategies

DOUGLAS T. BREEDEN*

This article presents a straightforward technique for computing solutions to discrete, multi-period consumption/investment problems. It solves for the optimal stochastic consumption plans, as well as the optimal dynamic trading strategies that maximize utility for an individual. The technique permits general utility functions that may or may not be time-separable. It also allows general changes in the investment opportunity set and allows the user to impose upper and lower bounds on trading behaviour. Divergent borrowing and lending rates can be handled, as can stochastic labour income risks. Computed solutions verify the predictions of well-known intertemporal works by Merton, Breeden and others.

(J.E.L.: G13).

1. Introduction

Over thirty years ago, Samuelson (1969), Merton (1969, 1971, 1973), Fama (1970), Hakansson (1970), Rubinstein (1976), Breeden and Litzenberger (1978) and Breeden (1979, 1984) introduced finance theory to multiperiod or 'intertemporal' models of consumption and portfolio choice. Those works showed us that optimal consumption and portfolio rules in the multiperiod world we live in are significantly different from those derived in single-period models. Although theorists are surely unanimous on this point, very few finance students and portfolio managers are being taught how to compute optimal multiperiod portfolio and consumption strategies. The intent of this paper is to present a straightforward approach to computing optimal dynamic consumption and portfolio strategies for relatively general problems.

The reasons for the lack of usage of multiperiod financial models are not hard to find. It is nearly impossible to solve analytically for the actual optimal policies with realistic fluctuations in investment opportunities,

* Correspondence: Dean and William W. Priest Professor of Finance, Duke University, Fuqua School of Business, Durham, North Carolina 27708, USA, Email: *doug.breeden@duke.edu*.
This is a revision of a paper first written at Duke in 1987 and presented at many universities, but never published. I am very grateful to Campbell Harvey, John Sprow and the participants of seminars at Berkeley, Chicago, Duke, MIT, North Carolina, Ohio State and Stanford for their comments. An MIT dissertation by Hua He (1990) and He and Pierson (1991) built on some of the ideas first presented here.

© Banca Monte dei Paschi di Siena SpA, 2004. Published by Blackwell Publishing Ltd, 9600 Garsington Road, Oxford, OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

given typical limitations on feasible trading strategies. Explicit solutions by academia's finest researchers for optimal dynamic trading strategies usually assume hyperbolic absolute risk aversion (HARA) utility functions and at most a single random element in investment opportunities.[1] Even within the HARA class of preferences and 1-variable opportunity sets, none of these researchers find explicit solutions that produce borrowing and lending strategies that take account of divergent borrowing and lending rates. Furthermore, meaningful income risk is typically not utilized in finding optimal portfolios.

The lack of closed-form analytical solutions is not an insurmountable barrier to computation of optimal trading and consumption strategies. The technique of stochastic dynamic programming can be used to find these optimal strategies for very general problems. Most theorists implicitly or explicitly assume that dynamic programming is the approach used to compute solutions. Thus, the optimization technique presented here should be viewed as an alternative to dynamic programming as a solution technique for the multiperiod consumption and investment problem.

If the dynamic programming technique were easy to use, this paper probably would not have been written. Unfortunately, it is not easy to apply stochastic dynamic programming to many interesting finance problems. One must first lay out the tree structure of events. Then, starting at the last period, T, in every possible state of the world one solves for the optimal consumption and portfolio rules for what will then be a single-period problem. Since when these solutions are computed, the optimal wealths that will be carried into time T in various states are unknown, solutions must be computed for a grid of possible levels of wealth in each state. Given those solutions for time T, one backs up to time $T-1$ and solves in each state of the world for optimal consumption and portfolio rules. Again, since it is unknown how much wealth will be carried into time $T-1$, one must solve for policies as a function of wealth at $T-1$. As the process computes backward through time-states, it computes the utilities of all possible wealth levels, all but one of which are subsequently found to be non-optimal. This is a computationally burdensome process.

The optimization method proposed here is straightforward, intuitive and forward-looking. The technique can be easily taught and programmed in worksheets and in relatively easy computer programs.[2] In the examples in the paper, one sees that all of the principles of intertemporal portfolio

[1] For examples, see Samuelson (1969), Merton (1969, 1971), Hakansson (1970), Kraus and Litzenberger (1975) Rubinstein (1976, 1981), Brennan (1979), Sundaresan (1984) and Cox *et al.* (1985a and b) and Cox and Huang (1987). Multiple state variables can be handled with logarithmic utility, since a log utility person has no net hedging demands.

[2] A simple program that solves the problems examined in this paper is available on request. Campbell Harvey at Duke has also programmed this method in GAUSS and in SAS.

© Banca Monte dei Paschi di Siena SpA, 2004.

theory are verified computationally, as, of course, they surely must be, if the theory is correct.

There are two key elements to the approach developed here. The first element is the exact representation of the pay-offs from dynamic trading strategies under uncertainty. An intuitive pay-off matrix much like that in Ross (1976) is developed to handle general problems in modeling dynamic trading strategies. Of course, this approach has as its foundations the scenario or 'time-state preference' approach of Arrow (1964), Debreu (1959) and Hirshleifer (1970). This type of model has been used extensively in theory by many authors and in practical applications by Cox *et al.* (1979) and Breeden and Litzenberger (1978).

The second important element of this paper is the multiperiod adaptation of a very insightful general optimization technique derived by William F. Sharpe (1978) and used by him in single-period portfolio problems.[3] A virtue of Sharpe's optimization method is the forward-looking nature of the technique, as opposed to the usual backward method of dynamic programming. This more intuitive approach to multiperiod strategies should enhance the understanding and usage of multiperiod planning models by portfolio managers.

In the remainder of the paper, Section 2 briefly presents Sharpe's portfolio optimization technique, and Section 3 gives the mathematical representation of dynamic trading strategies. Section 4 extends Sharpe's optimization technique to the intertemporal consumption-investment problem. Section 5 illustrates the technique and compares the computed optimal consumption and portfolio rules with the predictions of the theorists. Section 6 discusses limitations and extensions of the method, and Section 7 concludes the paper.

2. Sharpe's Portfolio Optimization Technique

Since Sharpe's optimization technique is not extensively used by academics, the general method and the principal application to finding mean-variance efficient portfolios will be sketched. Sharpe's principal observation is that an optimal portfolio is one that has equal expected marginal utilities of $1 invested in every asset. This follows from the first-order condition for an optimal portfolio, and is a well-known result (as Sharpe notes). The beauty of the technique is its economic intuition and its surprising computational efficiency. In the optimization literature, it is standard gradient computation method.

[3] More readable treatments with examples of this optimization technique are in Sharpe's 1987 book *Asset Allocation Tools*, and his popular textbook on *Investments*, 1985, Chapter 20.

© Banca Monte dei Paschi di Siena SpA, 2004.

To find mean-variance efficient portfolios, Sharpe computes the utility of a portfolio as its expected return less the variance of the portfolio divided by the risk tolerance of the investor. Thus, more risk-tolerant investors penalize risk less. Mathematically:

$$(1) \qquad \text{(Sharpe single-period)} \quad U = E(\check{r}_p) - Var(\check{r}_p)/T$$

where the symbols are standard for portfolio p. The marginal utility of asset i is given by differentiating (1) with respect to its portfolio weight:

$$(2) \qquad \text{(Sharpe single-period)} \quad MU_i = E(\check{r}_i) - (2/T)\, Cov(\check{r}_i, \check{r}_p)$$

As an asset's weight is increased in the portfolio, its marginal utility decreases in (2), since its covariance with the portfolio increases.

Sharpe's method of finding the optimal portfolio relies on the diminishing marginal utility of increasingly large holdings of a security. First, start with an arbitrary portfolio. For that portfolio, the vector of marginal utilities for all assets is computed, as in (2). Since the initial portfolio was arbitrary and probably not optimal, the marginal utilities for various assets are likely to differ. Then one does a swap of funds invested from the lowest marginal utility asset to the highest marginal utility asset. The gain in utility depends upon the size of the swap, so Sharpe derives the optimal size of the 2-asset swap.[4] The optimally sized swap is then made and the marginal utilities for all assets are recomputed based upon the new portfolio. A new swap is identified and made, marginal utilities recomputed, and so on. The process ends when all marginal utilities are essentially identical (to the degree of computational accuracy desired). At that point, since the optimal portfolio is unique if the covariance matrix is nonsingular, one knows that the optimal portfolio has been found.

One reason *not* to use this technique in an unconstrained single-period portfolio problem is that the optimal portfolio can be computed directly in one step for that case. To find the optimal portfolio, one needs only to invert the covariance matrix and multiply it times the vector of expected excess returns on assets. A difficulty in practice with the invert-and-multiply technique is that the portfolios found often violate maximum or minimum holdings constraints. For example, many individuals and institutions are prohibited from short-selling or have significant costs of short sales. If the direct solution suggests an infeasible negative position in such an asset, one cannot simply assume that the constrained optimum is to zero out that holding and hold all other assets in the same proportions. Asset returns' covariances make portfolio choices very interdependent. Thus, a quadratic programming approach with Kuhn–Tucker conditions is required for a

[4] The optimal swap size equals the marginal utility of the 'best' security minus the marginal utility of the 'worst', divided by the variance of their difference in returns. See Sharpe (1978, p. 11 or 1987, p. 59).

© Banca Monte dei Paschi di Siena SpA, 2004.

feasible solution to many realistic problems. In such situations, the power and intuition of this marginal utility-based algorithm are quickly appreciated.

Sharpe shows that his method easily accomodates upper and lower bounds on assets' holdings, and those bounds may be different for different assets. A simple check of each proposed swap to eliminate moves to infeasible holdings ultimately results in an optimal constrained portfolio. The Kuhn–Tucker conditions tell us that an optimal constrained portfolio has marginal utilities equal for all assets with interior solutions. Assets constrained by their upper bounds on holdings have marginal utilities greater than or equal to those of all interior holdings. Assets at their lower bounds on holdings have marginal utilities lower than or equal to those of all interior holdings.

The next section elegantly and precisely represents dynamic consumption and portfolio strategies that fulfill time-state by time-state budget constraints. Given this representation, it is shown that this section's single-period optimization technique can be used to obtain solutions to quite general multiperiod problems.

3. An Exact Representation of Dynamic Trading Strategies

The representation of time and uncertainty used is that of a tree structure of events. Once the tree structure of events is laid out, the returns from investments and from trading strategies are summarized by a **pay-off matrix, X**, similar to that of Nielsen (1974) and Ross (1976). However, the pay-off matrix used here explicitly considers contingent dynamic trading strategies. It is not assumed here that markets are complete or effectively complete or even Pareto optimal, so investors may not be able to insure against all risks that they wish to insure.

To illustrate this process, consider a simple trading problem with 3 dates, {0, 1, 2} and two assets – a stock portfolio (with prices like the Standard and Poor 500 Stock Index's values) and a real riskless bond that earns r% for the next period. At each date, there are four possible scenarios for current and prospective future returns, and these may be state-dependent. The four scenarios at each point in time are generated by letting the stock portfolio go up or down and by letting the riskless interest rate go up or down, with all combinations of stock and rate moves being possible. Let us assume that the tree of stock and bond returns is as in Figure 1, (where the date and state number are in parentheses).

Note that since the world ends for our trader at time 2, the interest rate branchings at time 2 for investments to time 3 are suppressed. No assumption about probabilities has been made. However, if probabilities of an up move are the same at each node, then the investment opportunity set in time-state (1,2) strictly dominates that in time-state (1,1), assuming

© Banca Monte dei Paschi di Siena SpA, 2004.

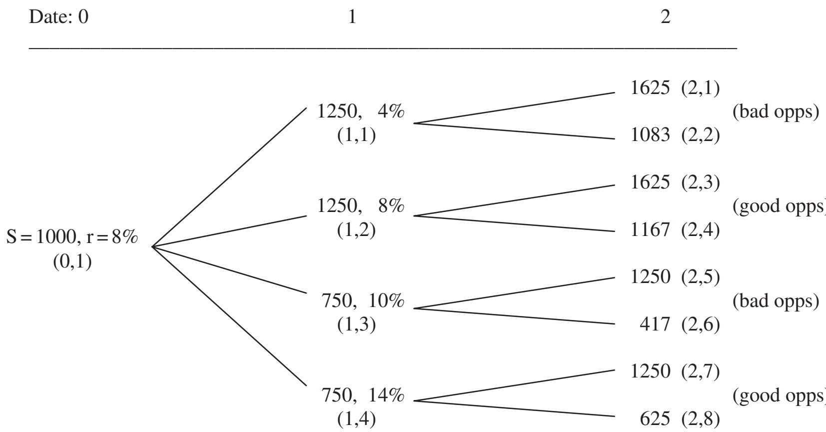

Figure 1: Tree of Stock and Bond Returns

that investors may not short securities. The dominance occurs since the riskless return is higher and the stock market return in (1,2) stochastically dominates that in (1,1). Similarly, the investment opportunity set in time-state (1,4) strictly dominates that in time-state (1,3). Thus, states (1,1) and (1,3) may be said to have 'bad investment opportunity sets', whereas states (1,2) and (1,4) have 'good investment opportunity sets'. This setup is useful in Section 6 in illustrating the principles of intertemporal portfolio theory of Merton (1973) and Breeden (1979, 1984).

Because there are 1, 4, and 8 states at dates 0, 1, and 2, respectively, there are a total of S = 13 'time-states', i.e., nodes in the tree. Each security or trading strategy's pay-offs on various dates and under various contingencies are fully described by the (positive or negative) cash flows that occur in those 13 time-states. The pay-off matrix that describes all possible trading strategies has 13 rows that represent the various time-states. Each possible trade's cash flows is a column in that matrix. In this problem, there are N = 10 state-contingent trades under consideration – buy stock or buy bonds in each of the 5 time-states at time 0 and time 1, {(0,1), (1,1), (1,2), (1,3), (1,4)}, and sell them the following period. Of course, stocks or bonds can be held for two periods by buying them at time 0 and again at time 1. Also, note that multiperiod investments that cannot be liquidated in intermediate periods are easily represented by columns in the pay-off matrix. The 13×10 pay-off matrix, **X**, is in Table 1.

The pay-off matrix gives the investments and pay-offs on a per share or per contract basis. Note that futures contracts, which require no initial investment, easily fit into this framework. The weight vector has as many elements as there are potential trades, e.g., 10 in the example above. For any trade for which there is a natural initial investment, it is convenient

© Banca Monte dei Paschi di Siena SpA, 2004.

Table 1: Pay-off Matrix X (S × N): Contingent Trades

Time-State		1 Stock 0,1	2 Bond 0,1	3 Stock 1,1	4 Bond 1,1	5 Stock 1,2	6 Bond 1,2	7 Stock 1,3	8 Bond 1,3	9 Stock 1,4	10 Bond 1,4
1	0,1	−1000	−100	0	0	0	0	0	0	0	0
2	1,1	1250	108	−1250	−100	0	0	0	0	0	0
3	1,2	1250	108	0	0	−1250	−100	0	0	0	0
4	1,3	750	108	0	0	0	0	−750	−100	0	0
5	1,4	750	108	0	0	0	0	0	0	−750	−100
6	2,1	0	0	1625	104	0	0	0	0	0	0
7	2,2	0	0	1083	104	0	0	0	0	0	0
8	2,3	0	0	0	0	1625	108	0	0	0	0
9	2,4	0	0	0	0	1167	108	0	0	0	0
10	2,5	0	0	0	0	0	0	1250	110	0	0
11	2,6	0	0	0	0	0	0	417	110	0	0
12	2,7	0	0	0	0	0	0	0	0	1250	114
13	2,8	0	0	0	0	0	0	0	0	625	114

(but certainly not necessary) for that trade's column to be divided by the cost of the investment. In Table 1, this would divide each column by the absolute value of its one negative amount; this puts the pay-off matrix in rate-of-return form for those assets. If that is done, then the weight vector gives the contingent dollar amounts invested in each such trade.

Complicated strategies and portfolios of assets are represented by the portfolio weight vector, **w**, which gives the number of shares or the sizes of the trades made. The net pay-offs (or costs if negative) in the various time-states from the trading strategy are given simply by the matrix product, **Xw**, which exists since **X** has dimensions S × N and **w** is N × 1.

To more fully see that this is an exact and very general representation of dynamic trading strategies, consider the following illustration related to the stock and bond example of Figure 1 and Table 1. Focus on Figure 1's tree and consider the asset allocation decision (stock-bond mix) at times 0 and 1. A portfolio manager has \$12,000,000 to invest initially. Assume that the initial allocation is 50% of funds in stocks and 50% in bonds, which implies that she buys \$6,000,000/1000 = 6000 'shares' of the S&P 500 stocks and buys \$6,000,000/100 = 60,000 bonds at time 0. Those positions are described by letting the first two elements of the portfolio weight vector be: $w_1 = 6000$ and $w_2 = 60{,}000$.

Next, assume that at time 1 our manager wishes to reduce her position in stocks if the market jumps up (taking profits), and to increase her stock position if the market drops down (taking advantage of the better values). The amount she wishes to shift to or from the bond market varies with the level of interest rates. The higher interest rates are, the more is shifted into bonds. An example of this dynamic trading strategy is in Table 2.

© Banca Monte dei Paschi di Siena SpA, 2004.

Table 2: Dynamic Trading Strategy Illustration

				Dollar Amounts		Shares Held	
t	Stock Price	Interest Rate	Wealth	Stocks	Bonds	Stocks	Bonds
0	1000	8%	12,000,000	$6,000,000	$6,000,000	6000	60,000
1	1250	4%	$13,980,000	4,500,000	9,480,000	3600	94,800
1	1250	8%	13,980,000	3,000,000	10,980,000	2400	109,800
1	750	10%	10,980,000	9,000,000	1,980,000	12000	19,800
1	750	14%	10,980,000	7,200,000	3,780,000	9600	37,800

Representation of this dynamic asset allocation strategy:

$w' = (6000,60000,3600,94800,2400,109800,12000,19800,9600,37800)$

This shows that with the proposed pay-off matrix for contingent trades, an exact representation of fairly complex trading strategies is easily done with the trade vector **w**. Thus, this generalized portfolio weight vector gives portfolio holdings for different assets at all dates and under all investment opportunity sets modeled.

Realistically for most individuals, wage income is a major source of funds. The level and uncertainty of future income importantly affects investment and (particularly) consumption decisions. An uncertain time- and state-dependent income stream is represented by the $S \times 1$ vector of dollar amounts, **y**. Correspondingly, the time and state-contingent consumption plan is represented by the $S \times 1$ vector, **c**. Without loss of generality, it is assumed that the first income element includes the trader's initial wealth, as well as any current income. The set of budget constraints for time-states is given by the vector equation:

$$\mathbf{c} = \mathbf{y} + \mathbf{X}\mathbf{w} \tag{3}$$

That is, at each date and state, consumption equals income plus net wealth withdrawals. Of course, alternatively, income equals consumption plus net investment. Note that (3) enforces a separate budget constraint for each time-state. This is in contrast to many formulations that have just one budget constraint that the present value of consumption equals the present value of income and wealth. If markets are not effectively complete, that simpler budget constraint may produce infeasible solutions. The approach here is more general.

The apparent simplicity of the vector of budget constraints, (3), may obscure the richness and generality of the formulation. Future income uncertainty and optimal consumption choices are being rather fully modeled – not just added on. To see this, think first of the income uncertainty already permitted in the tree of Figure 1 (which may be further 'branched' to provide a better representation of income uncertainty). If one works in the automobile industry, for example, then wage income is

© Banca Monte dei Paschi di Siena SpA, 2004.

probably tied somewhat to auto industry profits. Auto profits are dependent upon the general strength of the economy (which is reflected in stock prices), as well as upon the level of interest rates (since that affects the cost of financing durable goods). Thus, the income uncertainty of a manager for one of the auto companies might be modeled as in Table 3.

One can easily verify that the product of the pay-off matrix with the portfolio holdings vector, **Xw**, gives the net cash withdrawals from the portfolio's value. Negative numbers indicate investments made into the portfolio at the various dates with the various investment-income situations. Combining the net portfolio withdrawal amount (or saving) in a time and state with the income amount for that time and state gives the individual's funds for consumption in that time-state. This explains why (3) represents the state-by-state budget constraints.

Section 4 takes the uncertain income stream, **y**, and the intertemporal investment opportunity set, **X**, and searches for portfolio trading strategies, **w**, that result in the individual's optimal consumption plan, **c**. Specific formulae are derived for a relatively efficient modification of Sharpe's optimization technique to fit this multiperiod problem.

4. Optimal Dynamic Consumption and Trading Strategies: A New Approach

The optimal dynamic trading strategy is the one which maximizes the expected value of the individual's von Neumann–Morgenstern lifetime utility function for time-state consumption. That is, **w** is chosen to maximize U(**c**, $\boldsymbol{\pi}$), where the S × 1 vector $\boldsymbol{\pi}$ gives the probabilities of getting

Table 3: Uncertain Future Income Stream Illustration

t	S&P 500 Stock Index	Interest Rate	Initial Initial Wealth	Income Stream	Total Income Vector Representation **y**
0	1000	8%	$500,000	$100,000	$600,000
1	1250	4%		200,000	200,000
1	1250	8%		150,000	150,000
1	750	10%		85,000	85,000
1	750	14%		70,000	70,000
2	1625	4%		400,000	400,000
2	1083	6%		150,000	150,000
2	1625	8%		300,000	300,000
2	1167	4%		200,000	200,000
2	1250	10%		150,000	150,000
2	417	14%		50,000	50,000
2	1250	10%		125,000	125,000
2	625	14%		80,000	80,000

© Banca Monte dei Paschi di Siena SpA, 2004.

to the various time-states, given the initial conditions. For each date, the sum of the probabilities of various possible states equals unity. In this formulation the utility function for consumption may be time and state dependent. For the time-additive case, lifetime expected utility may be expressed as:

$$U(\mathbf{c}, \mathbf{f}) = \sum_t \sum_s \pi_{ts} u_{ts}(c_{ts}) = \boldsymbol{\pi}' \mathbf{u}(\mathbf{c}) \tag{4}$$

where $\boldsymbol{\pi}$ and $\mathbf{u}$ are $S \times 1$ vectors of stacked probabilities and utility functions for time-states, respectively, with rows corresponding to those in the pay-off matrix.

Application of Sharpe's optimization technique for the intertemporal problem is straightforward. Starting with Section 3's representation of feasible dynamic trading strategies, and given a vector of uncertain future incomes, the budget constraints of (3) are used to compute the consumption plan that follows from an initial (arbitrary) dynamic trading strategy, $\mathbf{w}$. From those contingent time-state consumption amounts, the direct utility function is used to compute marginal utilities for consumption in each time-state. Given those time-state marginal utilities, each asset's time-state pay-offs are used to compute the net marginal utility of another planned share purchase of the asset. Any contingent trade with a positive net marginal utility whose trade size can be increased (not against upper bounds on trade) is a candidate for a 'portfolio weight' increase. Similarly, any contingent trade with a negative net marginal utility whose trade size can be decreased (not against lower bounds on trade) is a candidate for a portfolio weight decrease.

Given these directions for portfolio improvement, a new trading strategy, $\mathbf{w}$, is considered. (The actual selection of the new trading strategy will be discussed later.) The same loop described is repeated. Consumption plans and marginal utilities are recomputed, followed by trades' marginal utilities, and a new trading strategy is selected. As in Section 2, the process ends when net marginal utilities for all trades are equal (to zero in this set-up), excepting trades with negative marginal utilities that cannot be decreased and trades with positive marginal utilities that cannot be increased.

Note that aside from possible upper and lower bounds, the choices of sizes of trades are unconstrained, since the costs of investment are already included as negative pay-offs. Increased funds required may be viewed as reducing consumption in the time-state when funds are required. Imposition of infinite marginal utility for zero or negative consumption makes it non-optimal to invest more wealth than one has in any time-state.

The principal choice available for minimizing computation time is in the updating of the trading strategy or portfolio vector, $\mathbf{w}$. Some updating

© Banca Monte dei Paschi di Siena SpA, 2004.

techniques are considerably more efficient than others.[5] Since PCs are widely used and the computation times involved are small enough, the hope of more widespread teaching and usage of multiperiod optimization models in finance seems a reasonable one.

Finding the most efficient change in a single trade's size requires consideration of the entire pay-off matrix, since pay-offs on different securities are correlated. Furthermore, in the intertemporal problem, many trades transfer wealth between the same two dates and states, which shows up in the pay-off matrix of Table 1 as two trade columns having zeros in the same rows of the pay-off matrix. For example, a riskless bond and a stock both held from time 0 to time 1 are clearly uncorrelated, but are significant substitutes, since both transfer wealth intertemporally over the same time period. Thus, in optimizing for the intertemporal problem, many securities that are not statistically highly correlated must be treated in computational design as if they are very highly correlated.

The problem with considering the entire pay-off matrix in choosing the optimal trade updates is that the pay-off matrix should be viewed as a huge matrix for many interesting problems. For example, having four branches at each node for ten time periods gives over a million possible states in the tenth period, since $4^{10} = 1{,}048{,}576$. Thus, it appears computationally prudent *not* to consider the entire pay-off matrix in selecting trade updates at each iteration.

If one starts simply by intending to change only one trade size at each iteration, which trade should be changed and by how much (ignoring the intercorrelations among trades)? To answer this, first one should look at the improvement in expected utility of lifetime consumption that a unit change in trade causes. This is given for each trade by its computed net (expected) marginal utility, given the current consumption plan. Trade i's net marginal utility is:

$$\mathrm{MU_i} = \sum\nolimits_t \sum\nolimits_s X^i_{ts} \pi_{ts} \acute{u}_{ts}(c_{ts}) \tag{5}$$

Let the time-states' probabilities multiplied by their marginal utilities for consumption (π_{ts} $\acute{u}_{ts}$) be stacked in a vector corresponding to the time-state rows of the pay-off matrix, and let that $S \times 1$ vector be denoted $\mathbf{U}_c$, which is a function of **c**. That vector will be called the vector of marginal utilities for consumption in the various time-states. Readers should understand

[5] For example, in an early application, optimal updating of trades one swap at a time reduced computation time by 90% over a naive, non-optimized updating method. That computation time was then reduced by another 75% with an optimal updating of two trades simultaneously at each iteration. Using GAUSS, the same technique solved the problem in one quarter of that time. The iterative process was stopped when the investor's shadow price for all trades was within .01% of zero. For example, this would happen when all investments that cost \$100.00 (in present value) had shadow values between \$99.99 and \$100.01. That tolerance can be made as small as is desired, but computation time gets larger.

© Banca Monte dei Paschi di Siena SpA, 2004.

that probabilities are included (which Hirshleifer (1970, Chapter 8) calls a 'util-prob' notation). Given that definition, the $1 \times N$ vector of trades' net marginal utilities, $\underline{\mathrm{MU}}$, can be expressed as:

$$\underline{\mathbf{MU}} = \mathbf{U_c X} \tag{6}$$

In choosing the optimal update, trades with positive net marginal utilities that are at their upper bounds should be ignored, since the direction of improvement is infeasible. Similarly, trades with negative marginal utilities that are at their lower bounds should be ignored. Of those trades remaining, the largest expected utility improvement per unit trade is given by the trade with the largest absolute value of net marginal utility. If that trade has a negative marginal utility, then its size should be reduced; if positive, the trade size should be increased.

To find the optimal size of the trade change, one can rely on Sharpe's point that changing the scale of a trade is subject to diminishing returns, due to the assumed diminishing marginal utility of consumption in each time-state. As a trade transfers consumption from one set of time-states to another set, those time-states with increased consumption have lowered marginal utilities, whereas time-states with decreased consumption have increased marginal utilities. Of course, these marginal utility changes are precisely related to the sizes and signs of the trade's pay-offs, so that the net marginal utility of the trade moves toward zero as the trade size is moved in the correct direction. Obviously, the rate at which marginal utility of a trade goes towards zero depends upon the rates of change of marginal utilities, i.e., the second derivatives of the utility functions.

The optimal size of a trade change is the size that (locally) drives the net marginal utility of the trade to zero on the next iteration, based upon the new consumption plan. Let dw^i be the change in the size of trade i, which may be positive or negative. For each time-state ts, the new consumption level as a result of pay-off from (or cost of) that trade change is:

$$c_{ts}^{new} = c_{ts}^{old} + x_{ts}^{i} dw^{i} \tag{7}$$

for all time-states ts, where the trade's pay-offs, $\{X_{ts}^i\}$, are positive in some time-states and negative in others. The new marginal utility of consumption in each time-state is approximately:

$$U_{ts}'^{\,new} = U_{ts}'^{\,old} + U_{ts}''^{\,old} x_{ts}^{i} dw^{i} \tag{8}$$

for all time-states ts. Based upon the new vector of marginal utilities for consumption, the trade's approximate marginal utility on the next iteration can be computed. The trade increment, dw^i, is chosen so that the new marginal utility is zero. Let the vector x^i be trade i's column of time-state pay-offs and let U_{cc} be the $S \times S$ diagonal matrix of second

© Banca Monte dei Paschi di Siena SpA, 2004.

partials (no cross-partials).[6] The new marginal utility for trade i is given by (6) and (8). Solving for zero:

$$\text{(9)} \qquad MU_i^{new} = x^{i\prime}[U_c^{old} + U_{cc}^{old} x^i (dw^i)] = 0$$

$$\text{(10)} \qquad \begin{aligned} &\Rightarrow dw^i = -x^{i\prime} U_c^{old} / [x^{i\prime} U_{cc}^{old} x^i] \\ &\Rightarrow dw^i = -MU_i^{old} / [x^{i\prime} U_{cc}^{old} x^i] \quad \text{(optimal 1-trade update)} \end{aligned}$$

Since U_{cc} is negative and the contributing x^i terms are squared in the denominator, the trade change has the same sign as the net marginal utility of the trade (as it should).

The optimal 1-trade update formula of (10) is directly analogous to Sharpe's optimal swap size (see footnote 4). The optimal trade increment equals the net marginal utility improvement divided by a squared pay-off term that is analogous to the variance of the swap.

The 1-trade updating approach works fine in solving the problem. However, in checking the paths of these calculations, it is clear that 'overshooting and reversals' occur that are not very efficient. For example, imagine what happens with a shift of funds at time zero from a bond to a stock. First the bond holding might be reduced, which transfers consumption intertemporally from time 1 to time 0. Then there is too much consumption at time 0, so stocks are purchased to push consumption back out to time 1, thereby effecting the switch from bonds to stocks at time 0. It took 2 steps to accomplish what could be done in 1. Furthermore, as one might imagine in a problem that acts as if it has a very high degree of multi-collinearity across trades, there is a great amount of shifting of weights back and forth among a few assets. So in fact it often might take 10 steps {trades 2,1,2,1,2,1,2,1,2,1} to make a switch of funds from trade 2 to trade 1. For that reason, it is important to derive the optimal adjustments in two trades simultaneously.

The optimal updates of two trades' sizes are found with a 3-step procedure just like that for 1-trade updates. First, compute the new consumption plan and time-states' approximate new marginal utilities for consumption in terms of the two unknowns that represent the sizes of the new planned trades. Secondly, compute what the new marginal utilities of trades will then be. Thirdly, solve for the trade sizes that give new marginal utilities of zero for the two trades. This can be done uniquely

[6] Actually, cross-partials should be in the matrix if time-complementarity of consumption is nonzero, as in time-multiplicative utility models. However, the number of calculations of second partials that are required at each iteration would be S^2, instead of S, which is a major drawback against a full matrix. With time-additive utility functions, all of those cross-partials are zeroes anyway.

© Banca Monte dei Paschi di Siena SpA, 2004.

by inversion of a 2×2 matrix, since there are two equations that are linear in the two unknown trade sizes.

With changes in two contingent trades, x^i and x^j, the new time-state consumption plan and marginal utilities of consumption are:

$$(7') \qquad c_{ts}^{new} = c_{ts}^{old} + x_{ts}^{i} dw^{i} + x_{ts}^{j} dw^{j}$$

for all time-states ts,

$$(8') \qquad U_{ts}'^{\,new} = U_{ts}'^{\,old} + U_{ts}''^{\,old}[x_{ts}^{i} dw^{i} + x_{ts}^{j} dw^{j}]$$

for all time-states ts.

The new marginal utilities for trades are then:

$$(9') \qquad \begin{aligned} MU_i^{new} &= x^{i\prime}[U_c^{old} + U_{cc}^{old}(x^i dw^i + x^j dw^j)] = 0 \\ MU_j^{new} &= x^{j\prime}[U_c^{old} + U_{cc}^{old}(x^i dw^i + x^j dw^j)] = 0 \end{aligned}$$

Optimality conditions for the optimal trade increments, $[dw^i, dw^j\}$, gives a 2×2 system of linear equations:

$$(11) \qquad \begin{aligned} (x^{i\prime} U_{cc}^{old} x^i) dw^i + (x^{i\prime} U_{cc}^{old} x^j) dw^j &= MU_i^{old} \\ (x^{j\prime} U_{cc}^{old} x^i) dw^i + (x^{j\prime} U_{cc}^{old} x^j) dw^j &= MU_j^{old} \end{aligned}$$

Letting **A** be the 2×2 matrix of coefficients in (11), the solution is:

$$(12) \qquad \mathbf{dw} = -\mathbf{A}^{-1}\ MU^{old} \quad \text{(optimal 2-trade and n-trade update)}$$

Upon reflection, one sees that (12) is also the general solution set for any number (N) of simultaneous trade updates, with **A** appropriately expanded to $N \times N$ and with MU^{old} expanded to $N \times 1$. Note that the diagonal of the coefficient matrix **A** is uniformly negative, so one's intuition that high marginal utility trades should be increased is the typical result of (12). However, large correlations with other trades updated can reverse the optimal direction, as one might expect.

5. Examples of Optimal Dynamic Trading Strategies

This section presents examples of solutions for optimal dynamic consumption and trading strategies. The illustrations considered here are chosen to illustrate the principles of intertemporal consumption and portfolio theory in this model. In particular, as described, optimal responses of individuals' portfolios to shifts in investment opportunities and to income risk are considered. Consumption policies are examined to verify the predictions of optimal responses for individuals with different degrees of relative risk aversion. Finally, an example is presented that shows sensible

© Banca Monte dei Paschi di Siena SpA, 2004.

portfolio strategies for the multiperiod borrowing and lending problem. Divergent borrowing and lending rates are examined, with prospective income streams and initial wealths that correspond to typical views of young and old investors.

For the first illustration, the contingent trade matrix used is in Table 4, which is identical to that described earlier in Figure 1 and Table 1, but is just normalized by putting pay-offs on a return-per-dollar-invested basis. By doing this, the optimal portfolios that are found (Table 5, described shortly) all represent the total funds (dollars) contingently invested in each trade. The *uncertain* income stream used is the one described earlier in Table 3 for the automobile firm manager. Income fluctuates substantially and directly with stock prices, and inversely with interest rates. A corresponding *certain* income stream, which pays the same expected income at each date as the uncertain stream, is also used. It is helpful in illustrating consumption and portfolio strategies' responses to more pure shifts in opportunities, without having the confounding wealth effects of a volatile income. Also, by comparison of optimal strategies with and without stochastic income, Mayers' (1972) optimal portfolio responses to 'non-marketable income' are verified.

In intertemporal problems, it is clear that individuals with different utility functions will have different consumption and portfolio strategies. Thus, it should not surprise anyone that utility functions must be specified to compute the optimal strategies. However, as noted earlier, the form of the utility function is essentially unrestricted, as long as equations for computing marginal utilities ($u_c > 0$) and their changes ($u_{cc} < 0$) can be written. Utility functions may change over time and in different states of the world. The specific utility functions examined here are pure power utility functions with constant relative risk aversion equal to RRA and with a pure rate of time preference equal to ρ. Mathematically, this utility function and its partials are:

$$
\begin{aligned}
& u = Ke^{\neg\rho t} c_{ts}^{1-RRA}, \text{ where } RRA >, K > 0 \text{ if } RRA < 1, K < 0 \text{ otherwise.} \\
(13) \quad & u_c = MU = K(1-RRA)\, e^{\neg\rho t} c_{ts}^{-RRA}, \\
& u_{cc} = -K(RRA)(1-RRA) e^{\neg\rho t} c_{ts}^{-RRA-1}
\end{aligned}
$$

These functions (multiplied by time-state probabilities) are used in the computer program to evaluate consumption plans at different dates. Equal probabilities at each branch are used in the illustrations, {1/4 for each state at $t = 1$, 1/8 for each state at $t = 2$}. A time preference parameter of $\rho = 3\%$ is used, with variations in that having predictable effects. A higher time preference parameter causes individuals to shift to more consumption earlier and less later, since ρ is a 'discount rate for future utility of consumption'.

© Banca Monte dei Paschi di Siena SpA, 2004.

Table 4: Uncertain Future Income Stream Illustration

	Income Paths		Trade Matrix										
State	Uncertain Income	Certain Income	Stock 0,1	Bond 0,1	Stock 1,1	Bond 1,1	Stock 1,2	Bond 1,2	Stock 1,3	Bond 1,3	Stock 1,4	Bond 1,4	
1	$600,000	$600,000	−1.00	−1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
1	$200,000	$126,250	1.25	1.08	−1.00	−1.00	0.00	0.00	0.00	0.00	0.00	0.00	Bad Opps
2	$150,000	$126,250	1.25	1.08	0.00	0.00	−1.00	−1.00	0.00	0.00	0.00	0.00	Good Opps
3	$85,000	$126,250	0.75	1.08	0.00	0.00	0.00	0.00	−1.00	−1.00	0.00	0.00	Bad Opps
4	$70,000	$126,250	0.75	1.08	0.00	0.00	0.00	0.00	0.00	0.00	−1.00	−1.00	Good Opps
1	$400,000	$181,875	0.00	0.00	1.30	1.04	0.00	0.00	0.00	0.00	0.00	0.00	
2	$150,000	$181,875	0.00	0.00	0.87	1.04	0.00	0.00	0.00	0.00	0.00	0.00	
3	$300,000	$181,875	0.00	0.00	0.00	0.00	1.30	1.08	0.00	0.00	0.00	0.00	
4	$200,000	$181,875	0.00	0.00	0.00	0.00	0.93	1.08	0.00	0.00	0.00	0.00	
5	$150,000	$181,875	0.00	0.00	0.00	0.00	0.00	0.00	1.67	1.10	0.00	0.00	
6	$50,000	$181,875	0.00	0.00	0.00	0.00	0.00	0.00	0.56	1.10	0.00	0.00	
7	$125,000	$181,875	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.67	1.14	
8	$80,000	$181,875	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.83	1.14	

© Banca Monte dei Paschi di Siena SpA, 2004.

Table 5: Optimal Dynamic Trading Strategies For Different Relative Risk Aversions

	Stock 0,1	Bond 0,1	Stock 1,1	Bond 1,1	Stock 1,2	Bond 1,2	Stock 1,3	Bond 1,3	Stock 1,4	Bond 1,4
Opportunity Set:										
Mean Return	0.00%	8.00%	8.33%	4.00%	11.67%	8.00%	11.11%	10.00%	25.00%	14.00%
Std. Dev'n	25.00%	0.00%	21.67%	0.00%	18.33%	0.00%	55.56%	0.00%	41.67%	0.00%
Reward/Risk	−0.32		0.20		0.20		0.02		0.26	
Uncertain Income:					(Dollars invested in each trade)					
RRA = 0.5	0	326,807	159,021	0	147,767	0	0	181,043	186,771	0
RRA = 2.0	0	308,971	0	152,937	0	121,145	0	162,110	34,899	113,679
RRA = 4.0	0	312,405	0	164,889	0	123,697	0	165,289	0	147,729
Certain Income:										
RRA = 0.5	0	319,976	153,159	0	161,105	0	25,699	133,189	182,512	0
RRA = 2.0	0	293,067	128,111	0	128,698	0	6,116	123,054	109,668	18,481
RRA = 4.0	0	289,395	73,748	51,965	88,305	36,409	2,873	121,926	53,098	69,728

© Banca Monte dei Paschi di Siena SpA, 2004.

Given these preferences and income and investment opportunities, a simple computer program that performs the optimization calculations in the last section was used to solve six problems. The six problems were generated by the two potential income paths in Table 4 (certain and uncertain), times three different levels of relative risk aversion {0.5, 2.0, and 4.0}. The optimal dynamic trading strategies for these 6 problems are in Table 5, and the optimal consumption plans are in Table 6. Bear in mind that these are only approximations to the optimal strategies, since iterations stopped when the shadow prices for all feasible trades were within 0.01% of zero.

Let us look first at the dynamic trading strategies of Table 5. A number of points are interesting and consistent with one's intuition. First, at time 0, all investors shun stocks entirely and buy bonds. That is not surprising, given that the mean return on stocks is 0 (8% less than bonds), yet stocks have a standard deviation of 25%, while the bonds are riskless. Note that states {(1,4), (1,1), and (1,2)} have the highest reward/risk ratios for stocks, and those are the states that have the stock/bond mix tilted most towards stocks. In both of the states with poor reward/risk ratios {(0,1),(1,3)}, stocks receive small or zero weights. Since uncertain income is tied somewhat to stocks' returns, income risk alone is like having a position in the stock market. Given this, it is not surprising that investors spend less on stocks when income is uncertain than when it is certain. In fact, with this uncertain income the very risk averse individual (RRA = 4.0) never buys stocks. Just as dramatically, in states {(1,1),(1,2)}, the stock-bond mix for the middle risk aversion person (RRA = 2.0) flips from 100% stocks to 100% bonds with the addition of this income risk. As expected, in all circumstances examined, the more risk tolerant individuals buy more stocks than do those with less tolerance for risk.

Next, let us look at the dynamic consumption plans (Table 6) and the consumption/investment mix. First, from Table 5, it is seen that all individuals invest more when income is uncertain than when it is certain. The greater their risk aversion, the greater the cutback in current consumption in the presence of such income risk. Individuals set aside more funds in investments to protect against potential adverse income fluctuations. In fact, this income risk shift is largely a wealth effect, (rather than an investment opportunity set effect), since individuals cannot invest more or less in this income (endowment) stream. From the left half of Table 6, it is seen that all individuals' optimal consumption plans here are highly correlated with their income streams. Thus, the shadow present value of the uncertain income stream is less than that of the riskless income stream, since it has the same expected payoffs at each date, but has positive consumption risk. With effectively less real wealth when income is risky like this, individuals consume less initially.

© Banca Monte dei Paschi di Siena SpA, 2004.

Table 6: Optimal Consumption Plans for Different Relative Risk Aversions

			Relative Risk Aversion				Relative Risk Aversion		
Date	State	Uncertain Income	0.5	2.0	4.0	Certain Income	0.5	2.0	4.0
0	1	$600,000	$273,193	$291,029	$287,595	$600,000	$280,024	$306,933	$310,605
1	1	$200,000	$393,931	$380,751	$372,509	$126,250	$318,666	$314,651	$313,083
1	2	$150,000	$355,185	$362,543	$363,700	$126,250	$310,720	$314,064	$314,082
1	3	$85,000	$256,908	$256,578	$257,108	$126,250	$312,936	$313,593	$313,997
1	4	$70,000	$236,180	$255,110	$259,669	$126,250	$289,313	$314,613	$315,970
2	1	$400,000	$606,727	$559,054	$571,484	$181,875	$380,931	$348,419	$331,791
2	2	$150,000	$287,818	$309,054	$321,484	$181,875	$314,613	$292,904	$299,834
2	3	$300,000	$492,097	$430,837	$433,593	$181,875	$391,311	$349,183	$335,994
2	4	$200,000	$337,916	$330,837	$333,593	$181,875	$332,240	$301,993	$303,615
2	5	$150,000	$349,147	$328,321	$331,818	$181,875	$371,216	$327,427	$320,782
2	6	$50,000	$249,147	$228,321	$231,818	$181,875	$342,661	$320,631	$317,590
2	7	$125,000	$436,286	$312,759	$293,411	$181,875	$486,061	$385,724	$349,861
2	8	$80,000	$235,643	$238,677	$248,411	$181,875	$333,968	$294,334	$305,613
Expected	t = 0	$600,000	$273,193	$291,029	$287,595	$600,000	$280,024	$306,933	$310,605
Expected	t = 1	$126,250	$310,551	$313,746	$313,246	$126,250	$307,909	$314,230	$314,283
Expected	t = 2	$181,875	$374,348	$342,233	$345,701	$181,875	$369,131	$327,577	$320,635
Avg Disc'td Consumption			$309,038	$305,935	$305,717	$297,934	$308,822	$306,792	$305,854
Std Devn	t = 1	$52,127	$65,942	$58,263	$54,954	$0	$12,560	$6,337	$1,051
Std Devn	t = 2	$108,913	$124,654	$100,592	$102,955	$0	$57,411	$51,365	$17,879

© Banca Monte dei Paschi di Siena SpA, 2004.

The more risk averse an individual is, the lower the mean and the lower the volatility of the consumption path, as shown by Breeden (1979, Sec. 5). This effect is verified here, as seen by comparing the means and standard deviations of consumption paths in Table 6.

As noted in Section 3 and is easily seen from Table 4, investment opportunities in state (1,2) dominate opportunities in state (1,1), and opportunities in state (1,4) dominate those in state (1,3). Breeden (1986, Sec. 6) showed that individuals more risk averse than the log (i.e., with RRA > 1) consume more currently when investment opportunities improve, *ceteris paribus*. This spreads the additional utility of good opportunities to consumption at each date. In contrast, individuals who are more risk tolerant than the log (i.e., those with RRA < 1), cut back current consumption and increase investment to take advantage of the better investment opportunities. The latter strategy leads to a higher mean lifetime consumption path and to a higher variance of consumption about that path. To see these effects, compare consumption in time-state (1,2) with that in time-state (1,1), and compare consumption in (1,4) with that in (1,3). To keep differential incomes from obscuring the effects, examine the panel where the income path is certain. Note that wealth is the same in each pair of states, since the S&P 500 is 1250 in (1,1) and (1,2) and is 750 in (1,3) and (1,4). The relatively risk tolerant individual (RRA $= 0.5$) does curtail consumption in the 'good opportunities' states and increases current consumption when the opportunities are poor. The very risk averse person (RRA $= 4.0$) does the opposite, smoothing the effects of the opportunity set into consumption at each date.

5.1. *Borrowing and Leverage Strategies*

One of the nicest features of Sharpe's optimization approach is the ease with which it handles limitations on short-sales, divergent borrowing and lending rates, and limitations on long positions. All of these are handled by the maximum and minimum quantities of trades that can be executed. Different trades have different maximums and minimums. By disallowing negative holdings of bonds and by representing borrowing by a positive early cash flow followed by negative later cash flows, the existence of both riskless borrowing and riskless lending at different rates can be modelled. As long as the borrowing rate is greater than or equal to the lending rate, there is no arbitrage. Of course, risky borrowing and risky lending can also be modeled.

Table 7 gives a new contingent trade matrix 'B', which represents a 3-date binomial trading set, $\{t = 0, 1, 2\}$. As a binomial (for simplicity), there are only 7 $(1 + 2 + 4)$ pay-off and investment nodes in the tree. At each node, one has three possible trades – borrowing, buying bonds, and buying stocks. Thus, with 3 trades at time 0 and at each of the two nodes at time 1, there are a total of 9 contingent trades to consider. To illustrate different

© Banca Monte dei Paschi di Siena SpA, 2004.

Table 7: Trade Matrix B

Date	State	Income Paths: Wealthy & Stable Income	Income Paths: Young MBA	Income Paths: Uncertain MBA	Stock 0.1	Bond 0.1	Borrow 0.1	Stock 1.1	Bond 1.1	Borrow 1.1	Stock 1.2	Bond 1.2	Borrow 1.2
0	1	$1,000,000	$50,000	$50,000	−1.00	−1.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00
1	1	$100,000	$100,000	$100,000	1.50	1.07	−1.11	−1.00	−1.00	1.00	0.00	0.00	0.00
1	2	$100,000	$100,000	$100,000	0.80	1.07	−1.11	0.00	0.00	0.00	−1.00	−1.00	1.00
2	1	$100,000	$200,000	$320,000	0.00	0.00	0.00	1.40	1.05	−1.09	0.00	0.00	0.00
2	2	$100,000	$200,000	$320,000	0.00	0.00	0.00	0.80	1.05	−1.09	0.00	0.00	0.00
2	3	$100,000	$200,000	$80,000	0.00	0.00	0.00	0.00	0.00	0.00	1.80	1.10	−1.15
2	4	$100,000	$200,000	$80,000	0.00	0.00	0.00	0.00	0.00	0.00	0.80	1.10	−1.15

© Banca Monte dei Paschi di Siena SpA, 2004.

supplies and demands for borrowing and lending, three different income paths are examined: (1) income for a wealthy ($W_0 = \$1$ million) individual with a fixed income of \$100,000 per year; (2) income for a 'typical young MBA with a certain income', who starts with an income of \$50,000, but *knows* that \$100,000 will be earned next year and \$200,000 in year 2; and (3) income for a 'young MBA with an uncertain income', who has year 2 income that has a 50–50 chance of either being \$320,000 or \$80,000 in year 2, and so has the same expected annual income as the certain MBA. The uncertain MBA's income is conditionally certain, in that if the stock market is up in year 1, then the \$320,000 income will be received in year 2. If the market is down in year 1, then the MBA knows that bad times are ahead and only \$80,000 will be earned in year 2. Also note that the investment opportunity set is especially favourable for stocks in state (1,2), dominating that in state (1,1) for stocks (although pure borrowers prefer state (1,1)).

Table 8 shows the optimal trading strategies for the borrowing/lending problem for individuals with different risk aversion and different income paths. As expected, wealthy individuals who invest money for later consumption tend to lend to the young MBAs. Furthermore, the more risk averse among the wealthy lend more risklessly than do the more risk tolerant wealthy. In fact, in this example, the most risk tolerant wealthy people ($RRA = 0.5$) do not buy any bonds – just stocks. The wealthy in this example never borrow. The more risk tolerant millionaires put more money in the stock market, whereas those who are very risk averse put more in Treasury bills. The young MBAs, who start with no wealth, consume more than they earn and also borrow to invest in stocks (a decadent group to be sure!). Even the young MBA who is most averse to risk ($RRA = 4.0$) borrows from a rising future income to fund current consumption, although very little is put into stocks by this person. Everyone responds to the outstanding investment opportunities in the stock market in state (1,2) by putting more funds into the market then than in state (1,1), despite their larger wealths in state (1,1).

Let us now compare the young MBA's strategy with the income uncertainty modelled with that of the certainty case. At time 1, the MBA knows what next year's income will be. If the stock market is high at time 1, the income next year will be high, whereas if the market is low at time 1, income at time 2 is low. One hopes to see and does see that borrowing is reduced in the situation where next year's income is sure to be low (1,2). One should not expect to reduce a low future consumption level further to increase a current normal consumption level. Almost all of the borrowing done in that low future income situation is used to finance investment in stocks, which present excellent, but risky prospects. Thus, the low income to be received at time 2 is used largely for time 2 consumption. On the other hand, when income is known to be high next year (1,1), the MBA

© Banca Monte dei Paschi di Siena SpA, 2004.

Table 8: Optimal Dynamic Trading Strategies With Leverage (RBorr > RLend)

	Stock 0,1	Bond 0,1	Borrow 0,1	Stock 1,1	Bond 1,1	Borrow 1,1	Stock 1,2	Bond 1,2	Borrow 1,2
Opportunity Set B:									
Mean Return	15.00%	7.00%	11.00%	10.00%	5.00%	9.00%	30.00%	10.00%	15.00%
Std. Dev'n	35.00%	0.00%	0.00%	30.00%	0.00%	0.00%	50.00%	0.00%	0.00%
Reward/Risk	0.23 & 0.11			0.17 & 0.03			0.40 & 0.30		
	(Dollars invested in each trade)								
Wealthy, Fixed Income									
RRA = 0.5	634,800	0	0	483,804	0	0	572,625	0	0
RRA = 2.0	295,830	287,497	0	137,903	229,999	0	176,443	84,830	0
RRA = 4.0	158,065	422,016	0	63,995	272,392	0	88,752	186,600	0
Young MBA (certain Y)									
RRA = 0.5	112,962	0	156,200	34,389	0	76,170	157,343	0	227,524
RRA = 2.0	27,330	0	84,806	6,656	0	78,320	35,516	0	113,466
RRA = 4.0	10,878	0	70,285	3,212	0	79,809	17,476	0	95,271
Young MBA (uncertain Y)									
RRA = 0.5	0	0	33,040	44,473	0	157,164	109,481	0	106,385
RRA = 2.0	0	0	36,422	10,297	0	132,416	22,749	0	30,860
RRA = 4.0	0	0	32,311	5,181	0	126,364	11,422	0	18,176

borrows against it to increase consumption at time 1, while maintaining a high consumption level at time 2 (an income smoothing strategy).

Clearly, many other interesting examples could be presented. However, this paper mainly serves to introduce a useful and theoretically sound optimization technique. The hope is that readers will explore many more examples. Section 6 examines some of the limitations of the technique and some extensions of it.

6. Limitations, Extensions and Embellishments

6.1. Utility Functions with Time Complementarity

There are a number of additional extensions of the model and examples of it that can fruitfully be examined. For example, a significant point is that utility functions with time complementarity can be used in this framework. The principal differences between time-multiplicative and time-additive utility functions are in the expressions for marginal utilities and second partials and cross-partials of utility with respect to time-state consumption. In the time-multiplicative case, the consumption level in one particular time-state affects the marginal utility and its changes in other time-states along the paths that go through the particular time-state examined. Also, in Equation

© Banca Monte dei Paschi di Siena SpA, 2004.

10, the matrix of second partial derivatives of utility with respect to time-state consumption is more full than the diagonal matrix of the time-additive case (see footnote 5). Still, the changes described are just in the formulae for computing values – the general optimization technique presented handles this case without modification.

6.2. Arbitrage Trading and Option Creation Strategies

Although it has not been illustrated, it should be clear that if there are arbitrage opportunities, they will be identified and exploited to the maximum degree possible. No matter what the utility function is, everyone likes a true arbitrage opportunity – no net investment, no risk of loss, non-negative returns in all time-states, and positive in at least one. Since an optimal *dynamic* trading strategy is found, even very complicated arbitrages are identified, possibly involving manoeuvres in and out of many different assets at different points in time. Similarly, 'quasi-arbitrage' opportunities, (which are assumed to be abnormally profitable strategies, given their very low risks), are likely to be optimal for most utility functions, and so are likely to be identified, almost regardless of the utility function used. A utility representation with very high risk aversion (such as a pure power function with a large negative exponent) can be used to identify only the lowest risk strategies.

6.3. The Major Limitation: Too Many Branches

The most obvious significant limitation of the general technique is the fact that the number of time-states and contingent trades can quickly become so huge as to overwhelm both computers and analysts. For example, binomial branching at each year for 30 years gives 1 billion (2^{30}) states in the 30th year! The nice, forward-looking tree structure that is such an intuitive advantage of this way of modeling becomes computationally infeasible with such huge trees. Still, probably the largest benefits of laying out dynamic consumption and trading strategies are in the first few years, as our abilities are suspect for seeing much of the economic structure beyond that. Thus, the most productive way to use the technique is to just attempt to model the first few years really well, rather than to get bogged down with huge numbers of very speculative trading opportunities in the distant future.

In carefully modelling only the first few years of what is truly a 30–50-year consumption and investment problem, a terminal indirect utility function for wealth carried into years past the last date can sometimes be used. For some tractable utility functions (such as some in the HARA class) and probabilistic assumptions for distant returns, solutions can be

© Banca Monte dei Paschi di Siena SpA, 2004.

found. Without a terminal utility of wealth, the optimization program would spend all wealth on consumption only in the first few years and would ignore income that is likely to be earned beyond the final date of analysis. For older people such an error would lead the program to too much current consumption, whereas for young people the error would probably be too little recommended current consumption.

7. Conclusion

The insights that have been developed in the finance literature about optimal dynamic consumption and portfolio rules have not been used much in practice. The likely reason for this is the difficulty of computing solutions to these complicated optimization problems. This paper combined an optimization technique developed by Sharpe with a new, exact characterization of the pay offs from dynamic trading strategies to find a simple solution method for optimal dynamic consumption and portfolio strategies. To illustrate the potential for widespread use, the technique was programmed and used on a personal computer to solve some interesting consumption, investment and income problems.

A number of examples were examined to verify the predictions of the theory about the nature of optimal consumption and portfolio rules. Of course, since these theoretical results have been around for a few years, any errors would have already been discovered. Thus, it is not surprising that the calculations 'verify the predictions of the theory'. The calculations really are just calculations of the formulae developed by theorists. Still, it is very useful to find that intuitively plausible optimal strategies can be easily computed for many problems that would take a long time to solve with the usual dynamic programming methods.

The technique appears to be rather general, in that it handles such things as income risk, divergent borrowing and lending rates, minimum and maximum trade positions, utility functions with time-complementarity, and allows almost any set of probability beliefs and risk aversion. Arbitrage opportunities can also be identified with the technique.

© Banca Monte dei Paschi di Siena SpA, 2004.

REFERENCES

K. J. Arrow (1964), "The Role of Securities in the Optimal Allocation of Risk-bearing", *Review of Economic Studies*, 31, pp. 91–6.

D. T. Breeden (1979), "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities", *Journal of Financial Economics*, 7, pp. 265–96.

D. T. Breeden (1984), "Futures Markets and Commodity Options: Hedging and Optimality in Incomplete Markets", *Journal of Economic Theory*, 32, pp. 275–300.

D. T. Breeden (1986), "Consumption, Production, Inflation, and Interest Rates: A Synthesis", *Journal of Financial Economics*, 16, pp. 3–39.

D. T. Breeden - R. H. Litzenberger (1978), "Prices of State-contingent Claims Implicit in Option Prices", *Journal of Business*, 51, pp. 621–51.

M. J. Brennan (1979), "The Pricing of Contingent Claims in Discrete Time Models", *Journal of Finance*, 34, pp. 53–68.

J. C. Cox - Chi-Fu Huang (1987), "Optimal Consumption and Portfolio Policies when Asset Prices follow n Diffusion Process", working paper, Sloan School of Management, MIT.

J. C. Cox - J. E. Ingersoll - S. A. Ross (1985a), "A Theory of The Term Structure of Interest Rates", *Econometrica*, 53, pp. 385–407.

J. C. Cox - J. E. Ingersoll - S. A. Ross (1985b), "An Intertemporal General Equilibrium Model of Asset Prices", *Econometrica*, 53, pp. 385–407.

J. C. Cox - S. A. Ross - M. Rubinstein (1979), "Option Pricing: A Simplified Approach", *Journal of Financial Economics*, 7, pp. 229–64.

G. Debreu (1959), *Theory of Value*, New York, Wiley.

E. F. Fama (1970), "Multiperiod consumption-investment decisions", *American Economic Review*, 60, pp. 163–74.

F. Grauer - R. Litzenberger (1979), "The Pricing of Commodity Futures Contracts, Nominal Bonds and Other Risky Assets Under Commodity Price Uncertainty", *Journal of Finance*, 34, pp. 69–83.

N. H. Hakansson (1970), "Optimal Investment and Consumption Strategies Under Risk for a Class of Utility Functions", *Econometrica*, pp. 587–607.

H. He - N. D. Pearson (1991), "Consumption and Portfolio Policies with Incomplete Markets and Short Sales Constraints: The Finite Dimensional Case", *Journal of Mathematical Finance*, 1(3), July, pp. 1–10.

J. Hirshleifer (1970), *Investment, Interest and Capital*, Englewood Cliffs Prentice-Hall.

© Banca Monte dei Paschi di Siena SpA, 2004.

A. Kraus - R. H. Litzenberger (1975), "Market Equilibrium in a State Preference Model With Logarithmic Utility", *Journal of Finance*,

D. Mayers (1972), "Nonmarketable Assets and Capital Market Equilibrium Under Uncertainty", in M. C. Jensen (ed), *Studies in the Theory of Capital Markets*, (Praeger).

R. C. Merton (1969), "Lifetime Portfolio Selection Under Uncertainty: The Continuous-time Case", *Review of Economics and Statistics*, 51, pp. 247–57.

R. C. Merton (1971), "Optimum Consumption and Portfolio Rules in a Continuous-time Model", *Journal of Economic Theory*, pp. 373–413.

R. C. Merton (1973), "An Intertemporal Capital Asset Pricing Model", *Econometrica*, 41, pp. 867–87.

N. C. Nielsen (1974), "The Firm as an Intermediary Between Consumers and Production Functions Under Uncertainty", unpublished PhD Dissertation, Stanford University.

S. A. Ross (1976), "Options and efficiency", *Quarterly Journal of Economics*, 87, pp. 75–89.

M. Rubinstein (1976), "The Valuation of Uncertain Income Streams and the Pricing of Options", *Bell Journal of Economics and Management Science*, 7, pp. 407–25.

M. Rubinstein (1981), "A Discrete-Tim Synthesis of Financial Theory", *Research in Finance*, Volume 3.

P. A. Samuelson (1969), "Lifetime Portfolio Selection by Dynamic Stochastic Programming", *Review of Economics and Statistics*, pp. 239–46.

W. F. Sharpe (1970), *Portfolio Theory and Capital Markets*, McGraw-Hill.

W. F. Sharpe (1978), "An Algorithm for Portfolio Improvement", Stanford Graduate School of Business Research Paper No. 475, October.

W. F. Sharpe (1985), *Investments*, 3rd edition, Prentice Hall, NJ.

W. F. Sharpe (1987), *Asset Allocation Tools*, 2nd edition, Scientific Press.

S. M. Sundaresan (1984), "Consumption and equilibrium interest rates in stochastic production economies", *Journal of Finance*, 39, pp. 77–92.

© Banca Monte dei Paschi di Siena SpA, 2004.