

Retirement Financial Planning: A State/Preference Approach

William F. Sharpe¹
February, 2006

PRELIMINARY DRAFT, PLEASE DO NOT QUOTE OR CITE

Introduction

This paper provides a framework for analyzing a prototypical set of decisions concerning spending, investment and insurance made by or for an individual or family following retirement. While the structure is formal it is designed to be used eventually for actual decisions. Hence attention will be paid not only to the overall structure but also to the design of algorithms and approximation methods that can facilitate low-cost determination of desirable strategies.

We utilize the discrete-time, discrete-state approach of [Arrow 1953] and [Debreu 1959] as well as a number of the concepts in [Sharpe 2006]. In this paper we also assume that markets are sufficiently complete in the sense defined in the latter. Hopefully, subsequent papers will cover cases in which markets are incomplete and/or there are costs associated with some or all financial products.

The Retirement Financial Planning Problem

For convenience we will focus on a single individual whom we will call “the investor”. In practice decisions are often made for two or more people, acting as a family unit, however our approach can easily be adapted to handle such cases by expanding the number of personal states considered. We will not differentiate between cases in which the investor selects his or her own retirement financial plan and those (more likely) instances in which a professional advisor provides assistance.

I

¹ This paper has benefited greatly from comments and suggestions provided by Wei Hu, Jason Scott, Jim Shearer and John Watson of Financial Engines, Inc., Steve Grenadier of Stanford University and Geert Bekeart of Columbia University

The investor begins with a set of state-dependent “outside sources” of income such as social security, a defined benefit retirement plan, etc. He or she also has a given amount of *investable wealth* that can be used to provide additional spending in different states of the world. A retirement financial plan is a set of state-dependent amounts of spending with a present value equal to the initial investable wealth.

The investor’s goal is to select the most desirable feasible financial plan. We operationalize this concept with the notion of a measure associating a level of *desirability* with any financial plan. Desirability depends on the total amount to be spent in each state and the probability of each state. For simplicity we assume that the contribution of each state to total desirability can be computed separately, with the overall desirability of a plan equal to the sum of the desirabilities associated with the states.


As will be seen, our concept of desirability is formally identical to the economist’s standard construct of expected utility. However, we adopt a different nomenclature to reflect a salient characteristic of the retirement financial planning problem. There are states of the world for an individual (for example, 90 years old, in a skilled nursing home with Alzheimer’s disease) in which the concept of utility at the time is fanciful at best. For this reason it is best to interpret the utility associated with a state-contingent amount to be spent in such a state as the utility that the individual associates with knowing today that this will be the situation in that future state if it comes to pass. For this reason we term the probability-weighted sum of state-dependent “utilities” the investor’s assessment of the desirability (today) of the specific set of amounts to be spent in possible future states of the world.

Market and Personal States

A key aspect of the retirement financial planning problem is the identification and analysis of states of the world that differ in ways relevant for many people from those that differ in ways relevant for the investor in question. We illustrate this and other aspects of our approach with a simple two-date example.

I

There are two dates: Now and Next Year. There is one state this year. The market next year can be in one of three possible *market states*: Bear, Normal or Bull. The investor can be in one of two possible *personal states*: Alive or Dead. Thus there are 7 states in all. Diagrammatically:



Each x value indicates an amount available to be spent in the state in question. The total amount spent in state s , denoted by z_s will equal this amount (x_s) and the amount, if any, from outside sources which we will denote y_s :

$$z_s = x_s + y_s$$

The complete set of x values will be denoted X , the complete set of y by Y and the complete set of z values by Z . Y is taken as given; the only decision variables are the values in X .

The Optimal Retirement Financial Plan

The investor's preferences are reflected in a series of utility functions, one for each state. We denote the utility function for state s as u_s . The utility associated with total spending in state s is thus $u_s(z_s)$.

We impose some restrictions on the utility functions. For each one, marginal utility ($u'_s(z_s)$) must be continuous and decrease with total spending (z_s). As shown in [Sharpe 2006] this allows a wide range of representations of investor preferences, including quadratic utility, constant relative risk aversion utility, HARA utility and piecewise marginal utility functions reflecting some of the characteristics of behavioral models of investor preferences.

Importantly, the utility functions for all states that occur at the same date (year) and reflect the same personal state will be the same (for example, u_2 , u_3 , and u_4 in the example).

Associated with each state is a probability, π_s . The set of all such probabilities will be denoted Π . The probability for a state will equal the product of the probability of the associated market state times the associated personal state. The probability of the present state is, of course, 1. The sum of the probabilities of all possible states at a given date will also equal 1.

Associated with each state is a *state price*, p_s . This is an amount that can be paid today to provide 1 unit of spending in state s (but only in state s). We assume that the market is sufficiently complete that any desired amount x_s can be provided in state s for an expenditure from investable wealth equal to $p_s x_s$. The set of all state prices will be denoted P . The cost of a retirement plan X is thus:

$$c = \sum_s p_s x_s$$

The desirability of retirement plan X is given by:

$$d = \sum_s \pi_s u_s(z_s)$$

Or:

$$d = \sum_s \pi_s u_s(x_s + y_s)$$

Given an amount of investable wealth w , the goal is to select X to maximize d subject to the constraint that $c = w$. More simply put, we want to find the most desirable plan the investor can afford. We will term the solution to this problem the *optimal retirement financial plan*.

Credit Constraints

We allow negative x values representing agreements that require payment in some states of the world. This is feasible if the total consumption in each state (z_s) is positive, since the payment (x_s) can be made from income (y_s) in the state.


Rather than impose a constraint that no elements of Z can be negative we require utility functions for which marginal utility is infinite for zero consumption (or, possibly, a positive feasible level of consumption). In this way the marginal utility acts as a barrier or penalty function that guarantees that the optimal plan will not violate credit constraints in any state.

Efficient Plans

For a given cost c there will be an optimal plan X . We term such a plan an *efficient plan*, since it provides the maximum desirability (d) for the specified cost (c).

Now consider the set of efficient plans, one for each possible level of cost c . Given the restrictions we have imposed, there will be a monotonic and increasing relationship between d and c : more expensive efficient plans will be more desirable.

The figure below provides an hypothetical example.



In this case the investor's wealth equals c_a . The current actual (suboptimal) plan provides a desirability of d_a . However, an optimal plan could provide a desirability of d_e . It would be tempting to consider the ratio d_a/d_e or the difference d_e-d_a as a measure of the relative efficiency of the current plan. However, this would be based on the premise that the measure of desirability represents some fundamental aspect of the investor's happiness. This is far too much to assume. We use the concept only as a means to summarize the investor's willingness to substitute consumption in one state of the world for consumption in another. In this context, any linear transform of all the utility functions would lead to the same choice of a retirement financial plan. Thus we could multiply every utility function by one positive constant and add another constant to each function without changing the representation of the investor's preferences. However, the former would change the difference d_e-d_a and the latter the ratio d_a/d_e .

For this reason we choose monetary measures of the relative efficiency of a plan. We find the efficient plan (*) that provides the same level of desirability (d_a) as the actual plan, then determine its cost (c^*). The ratio of the cost of an equally desirable efficient plan to the cost of the actual plan (c^*/c_a) or the difference between the costs of the plans (c_a-c^*) can then serve as measures of relative efficiency. Clearly these will be unaffected by any linear transformation of the utility functions. They also have the advantage of being based on monetary measures and being easily understood.

Marginal Utility Functions

Key to choices made under uncertainty are the investor's *marginal utility functions* for each of the states. The marginal utility for consumption in a state is the change in utility per unit change in consumption:

$$m_s(z_s) \equiv du_s(z_s)/dz_s$$

In our setting the investor's decisions depend only on the marginal utility functions. Of course the marginal utility of consumption in a state will depend on the amount consumed in that state.

A common assumption is that the investor's preferences display constant relative risk aversion:

$$m_s = a_s z_s^{-b_s}$$

Such a function is infinite at $z_s = 0$ which meets our requirements for credit constraints. However, other marginal utility functions can be used including HARA functions of the form:

$$m_s = a_s (z_s - q_s)^{-b_s}$$

In this case, q_s , the minimum acceptable consumption in the state must be non-negative and only plans for which z_s exceeds q_s in each state can be considered. As can be seen, when $q_s=0$, this becomes a constant relative risk aversion function. For positive values of q_s , it exhibits decreasing relative risk aversion.

It is entirely possible to represent an investor's preferences for consumption in a given state with a marginal utility function composed of a series of pieces, one for each of several ranges of consumption in that state. We require only that marginal utility be infinite at a non-negative level of consumption, that marginal utility be continuous and that it decrease as consumption increases. These conditions can be satisfied by, for example a series of connected HARA functions which can reflect at least some of the characteristics found in experimental studies of choice under uncertainty, as discussed in [Sharpe 2006].

Finally, preferences in one state can be represented with functions that differ in type as well as in specifics (parameters) from those in another state at the same or at a different date.

Investor Reservation Prices

Given a plan, it is straightforward to compute an investor's *marginal rate of substitution* for consumption in one state vis-a-vis consumption in another. This is the rate at which he or she could substitute consumption in one state for the other without changing the overall desirability of the plan. Of particular interest is the rate at which the investor would substitute consumption in a future state for present consumption. We term this his or her *reservation price* for the state. Since the present is state 1, the reservation price for state s will equal:

$$\frac{\pi_s m_s(z_s)}{\pi_1 m_1(z_1)}$$

Of course the probability of the present equals 1 so only the marginal utility of present consumption need be included in the denominator.

Optimality Conditions

A plan will be inefficient if the investor's reservation price for any state differs from the price for that state. If the reservation price is higher, the investor can obtain a more desirable plan by increasing consumption in the future state and reducing present consumption. If the reservation price is lower, the investor can obtain a more desirable plan by decreasing consumption in the future state and increasing it in the present. A necessary condition for a plan to be efficient is for the investor's reservation price to equal the actual price for every state:

$$\frac{\pi_s m_s(z_s)}{m_1(z_1)} = p_s$$

These are often termed the *first-order conditions* for optimality. In our setting these conditions are both necessary and sufficient for a plan to be optimal.

Finding an Efficient Plan

Finding an efficient plan would appear to require solving a complex nonlinear optimization problem. In our setting, however, a much simpler procedure can be utilized.

Consider an efficient plan Z' . It will have cost c' and desirability d' . The first element z_1' will indicate the amount to be spent in the present period. Now consider the set of all efficient plans. Efficient plans with greater cost will have greater desirability. They will also provide more consumption in the present period.

Now assume that we know the amount of present consumption z_1' for an efficient plan. The amount to be consumed in state s can then be determined using the first order conditions. Rearranging the prior expression gives:

$$m_s(z_s') = m_1(z_1') p_s / \pi_s$$

This can be solved for z_s' (and hence x_s'). For example, with a HARA marginal utility function:

$$a_s (z_s' - q_s)^{-b_s} = m_1(z_1') p_s / \pi_s$$

Taking the logarithms of both sides provides a simple equation that can be solved for z'_s . Subtracting y_s gives x'_s .

With piecewise marginal utility functions a solution can be obtained for each marginal utility function, then the one with a value falling within the relevant range utilized.

Given z_1 , the set of first order conditions for all future states can be solved to find the other elements of the efficient plan with initial consumption equal to z_1 . The cost of the plan can then be computed. This provides one point on the efficient frontier shown in the earlier graph.

Finding the Optimal Plan

The optimal plan for the investor will be the efficient plan with a cost equal to his or her wealth. Since efficient plans can be computed quickly, it is computationally efficient to find as many such plans as required to determine the optimal plan. First an initial consumption z_1 is selected, the corresponding efficient plan determined and its cost computed. If the cost equals the investor's wealth, this is the optimal plan. Otherwise a different value is selected for z_1 and a new efficient plan and corresponding cost determined. The process is continued until an efficient plan with a cost equal to the investor's wealth is found.

In some special cases only one iteration will be required. For example, if all the marginal utility functions are of the HARA type with the same risk aversion coefficient (b_s), the cost of an efficient plan will be proportional to the initial consumption z_1 . Any arbitrary value of z_1 can be selected for the first solution and the corresponding cost computed. Then every value of z_s can be set to equal the ratio of the investor's wealth to the cost of the initial plan.

For cases with greater variation in marginal utility functions, efficient iteration methods can be used to home in on the optimal portfolio. To begin, a low value of z_1 is used to find a plan with a cost below the investor's wealth and then a high value of z_1 is used to find a plan with a cost above the investor's wealth. Next an intermediate value of z_1 is selected and a new plan determined. It then replaces either the low or high plan, depending on the relationship of its' cost to the investor's wealth. The process continues until a plan with a cost sufficiently close to the investor's wealth is found.

Parsimony

Clearly there are many possible market states, many personal states and, as a result, a great many specific states.

As a practical matter it is important to reduce the dimensionality of the retirement financial plan as much as possible. To see how this can be accomplished consider the first order condition for the optimal amount to be consumed in state s :

$$m_s(z'_s) = m_1(z'_1)p_s/\pi_s$$

Now, consider two states with the same marginal utility function and the same ratio of price to probability (termed price per chance in [Sharpe 2006]). Clearly the optimal plan will require the same level of consumption in each of the states. Suppose that we choose to combine the two states into a single new state. The new state will have a price equal to the sum of the prices of the initial states and a probability equal to the sum of the probabilities of the initial states. Moreover, it will have the same ratio of price to probability as each of the original states. An optimal plan determined using this new state instead of the two original states will correctly call for the same amount of consumption as was originally indicated for each of the component states. The amount to be spent from investable assets in each of the component states can then be determined by subtracting the outside income from the total amount to be spent.

This example provides the basis for aggregation of “sub-states” into a smaller number of more aggregate states. For purposes of finding the total amount to be consumed, two or more states can be combined if the states:

1. have the same marginal utility functions, and
2. have the same ratios of state price to state probability

Note that both these conditions must hold to allow states to be aggregated.

If market returns are unrelated to an investor’s marginal utility functions it may be possible to aggregate security market outcomes into a series of overall market states (the *superstates* proposed initially by [Hakansson 1976]), greatly reducing the dimensionality of the retirement financial planning problem. Moreover, it may be plausible to assume that the market is sufficiently complete if dynamic strategies can be used to obtain payments contingent on the realizations of such superstates, making the approach described here relevant.

As a practical matter, of course, there could still be far too many states for practical solution since the overall market portfolio's return can take on many different values. However, in many cases discrete approximations can be utilized effectively. For example, one could approximate a lognormal distribution of return and derive a set of state prices using the approach in [Sharpe 2001] or an alternative procedure. Importantly, this should suffice in cases in which the goal is to select a strategy that will provide payments related to overall market performance in some relatively simple manner.

Of course the solution to any retirement financial planning problem will be only as good as the assumptions about capital markets that underlie it. The better the assumptions about the determinants of asset prices, the better will be the recommended plan.

Incomplete Markets

Unfortunately, in many cases markets may not be sufficiently complete to allow the fully optimal solutions of the type derived here for many investors' retirement financial planning problems. However, a great deal can be learned about such strategies by assuming such perfect and frictionless conditions. Moreover, some aspects of incomplete markets can be accommodated by aggregating states for which detailed markets do not exist.

When there is no way to obtain different amounts of spending in states with different marginal utility functions, it may be possible to reformulate the problem to determine a plan that will be optimal, given such a constraint. Consider, for example, a case in which annuities are infeasible or unduly expensive so that the same amount must be spent in the states "alive" (a) and "dead" (d). The expected marginal utility for an expenditure of z' conditional on one of the two states being realized will thus be:

$$\pi_a m_a(z') + \pi_d m_d(z')$$

This can be regarded as the marginal utility function for a single aggregate state that combines the two personal states. Given such aggregation, a set of constrained efficient plans can be found. The efficient frontier will generally fall below that for unconstrained solutions, providing a lower desirability (d) for any given cost (c). The inefficiency of the resulting plans can provide a measure of the value of new personalized financial instruments that could help complete the markets.

Further Research and Application

This paper is intended to provide a framework for examining aspects of retirement financial planning. At best it is only a first step towards helping individuals make efficient decisions concerning this stage of their financial and personal lives. Hopefully, much more research will build on this base leading in time to applications of significant practical value.

References

Arrow, Kenneth J., 1953, "Le Rôle de valeurs boursières pour la répartition le meilleure des risques," *Econométrie, Colloques Internationaux du Centre National de la Recherche Scientifique* 11, pp.41-47.

Debreu, Gerard, 1959, *The Theory of Value*, Wiley and Sons, New York.

Hakansson, Nils, 1976, "The Purchasing Power Fund: A New Kind of Financial Intermediary," *Financial Analysts Journal*, December 1976.

Sharpe, William F., 2001, "Individual Risk and Return Preferences: A Preliminary Survey," September, 2001, www.stanford.edu/~wfsharp/art/rrsurvey/vienna2001.htm

Sharpe, William F., 2006, *Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice*, Princeton University Press, November.