The Distribution Builder:
A Tool for Inferring Investor Preferences

William F. Sharpe, Daniel G. Goldstein and Philip W. Blythe*

October 10, 2000

Abstract

This paper describes *the Distribution Builder*, an interactive tool that can elicit information about an investor's preferences. Such information can, in turn, be used when making decisions about investment alternatives over time for that investor. The approach can also be employed when conducting surveys designed to obtain data on the cross-section of investor preferences. Hopefully, such data can provide insights that can lead to more realistic models of equilibrium in capital markets.

The approach asks an investor to choose among alternative probability distributions for end-of-period wealth, where only distributions with similar overall costs are allowed. Importantly, the cost of any distribution is consistent with a model of equilibrium pricing in capital markets. We show how such a model can be calibrated and how information about an investor's marginal utility of wealth can be inferred from his or her choice of a distribution.

Introduction

Models in Financial Economics are frequently built on assumptions about the preferences of investors. For example, the original Capital Asset Pricing Model[1] followed the approach developed by Markowitz[2], in which each investor is assumed to wish to maximize a linear function of the mean and variance of portfolio return. Moreover, investors are assumed to differ in their willingness to substitute mean return for variance of return. Given a world of such investors, the CAPM derives equilibrium conditions for security prices and the relationships among risks, correlations and returns. Its key implications for optimal portfolio holdings are that the (wealth-weighted) average investor should hold all securities (the market portfolio) while investors with less (more) tolerance for risk should invest less (more) in the market portfolio and more (less) in a riskless security.

Other equilibrium asset pricing models use more detailed models of investor preferences. Many start from explicit assumptions about the relationship between an investor's utility and wealth or consumption. For example, a multi-period model of equilibrium may characterize an investor's preferences in a manner that involves both a measure of risk tolerance and another relating to time preference[3]. As with simpler models, in equilibrium the wealth-weighted average investor should hold the market portfolio while others should adopt different strategies, depending on their relative degrees of risk tolerance and time-preference. More recent models utilize utility functions with more parameters and hence obtain results that imply more diversity in optimal portfolio holdings[4].

Most asset pricing models focus on the prices of assets in equilibrium and the resulting relationships among risks, returns and correlations of returns. For such purposes the use of relatively simple characterizations of investor preferences may be perfectly reasonable. However, to explain actual investor holdings or to advise investors concerning optimal strategies it may be necessary to adopt a richer characterization of investor preferences or, at the very least, to have a better understanding of actual preferences so that a parsimonious characterization of such preferences can be utilized.

Not surprisingly, when choosing a form of utility function theorists have taken into account not only plausibility but also analytic tractability. This has led to "traditional assumptions" in one area that differ from those in another. For example, many life-cycle models of investor behavior assume that investor preferences exhibit constant relative risk aversion[5]. On the other hand, many models that focus on information assume that investor preferences have constant absolute risk aversion[6]. Any given investor could have one or the other, a function that exhibits one kind of behavior in one range of outcomes and the other in another range, or an entirely different type of function. But it cannot be the case that every investor has both functions at once.

To understand at least some phenomena and to offer investors the best possible advice, it is desirable to know more about the actual preferences of individuals. There are two ways to approach this subject. The first, common in the finance literature, is to make assumptions about preferences, imply equilibrium implications, then evaluate the degree of consistency of the implications with empirical data[7]. The second, common in the literature in cognitive psychology and behavioral finance, is to present subjects with alternative choices and infer preferences from the resulting selections. Prominent in the latter tradition is the work of Kahneman and Tversky[8], which showed that individuals in experimental settings make choices that are inconsistent with some of the standard properties of the utility functions and axioms of choice used in most theories in Financial Economics.

One of the key findings from the psychological studies of choice under uncertainty is the importance of *framing*. Subjects presented with alternatives that are the same in objective terms will often make different selections if the alternatives are described in one way rather than another[9]. This makes it imperative that attempts to elicit an investor's true preferences involve choices among alternative outcomes that are as similar as possible to those available in actual capital markets, with the alternatives stated in terms that are relevant for the individual in question.

This paper describes a method designed to aid in this process. We introduce a tool called the *distribution builder* that allows an individual to examine different probability distributions of future wealth and choose a preferred distribution from among all alternatives with equal cost. An important features is the requirement that, in order to make the choice set realistic, the cost of each distribution is consistent with an equilibrium model of asset pricing in capital markets . Finally, the nature of the distribution is presented in a manner designed to be easily understood by those not familiar with probabilistic analyses.

We envision two major uses for this tool. The first is normative in nature. Once an individual has chosen a distribution, it is possible to determine an investment strategy through time that will provide that distribution. With this information a third party could provide advice or implementation to help the investor meet his or her goals.

The second application relates to positive models of asset pricing and investor behavior. Once an individual has chosen a distribution using the tool it is possible to make inferences concerning his or her utility function. Given experimental data of this type from a number of individuals it should be possible to better select a set of parsimonious assumptions about investor preferences for building equilibrium capital market models.

As an illustration of these two types of application, consider investments such as equity index-linked notes that offer "downside protection" and "upside potential". An investor who purchases such an instrument may not fully understand the trade-offs involved in choosing the associated distribution over one that would result from a more traditional strategy such as a combination of an equity index fund and a riskless asset. The distribution builder can help make such trade-offs clear and allow an investor to make a more informed choice among alternative strategies. Turning to considerations of equilibrium we know that in order for markets to clear, a minority of investors should adopt such strategies with an equal minority (in value terms) adopting strategies with the opposite characteristics[10]. However, theory alone cannot provide information concerning the sizes of such minorities. In equilibrium, when investors fully understand the trade-offs, should 45% purchase downside protection, 45% provide it and only 10% adopt more traditional investment strategies, or are the percentages 1%, 1% and 98%, or even 0%, 0% and 100% (as in models such as the CAPM)? The answer will ultimately depend on the cross-sectional distribution of investor preferences. Widespread experimentation with tools such as the distribution builder should make it possible to better assess the characteristics of investors in a given market.

The plan of the paper is as follows.

Section 1 shows how a distribution builder presents a probability distribution in terms easily understood and manipulated by users. It also describes the role of the budget constraint in limiting possible choices. Section 2 describes a method used to compute the least cost of a distribution, given a set of Arrow-Debreu prices for possible future states of the world, where each state is equally probable. Section 3shows how attributes of a user's utility function can be inferred from his or her choice of distribution, given the underlying Arrow-Debreu prices. Section 4 shows how a simple binomial pricing model can be used both to compute the required Arrow-Debreu prices and to determine a specific dynamic strategy that will provide a chosen distribution. Section 5 provides a summary and conclusions as well as suggestions for further research.

1. The User Interface

A Distribution Builder lets people build and explore different probability distributions of a future source of utility, such as wealth or retirement income, under the constraints of a fixed budget. Figure 1 shows a typical user interface. The main parts of the tool are the large square playing area, a given number of "people" (here, 64), the reserve row (along the bottom of the playing area), and the *budget meter*. In this case the source of utility is income per year after retirement, expressed as a percentage of income in the year prior to retirement. Here, the user is told that the tool can help make decisions about the likely ranges of retirement income.

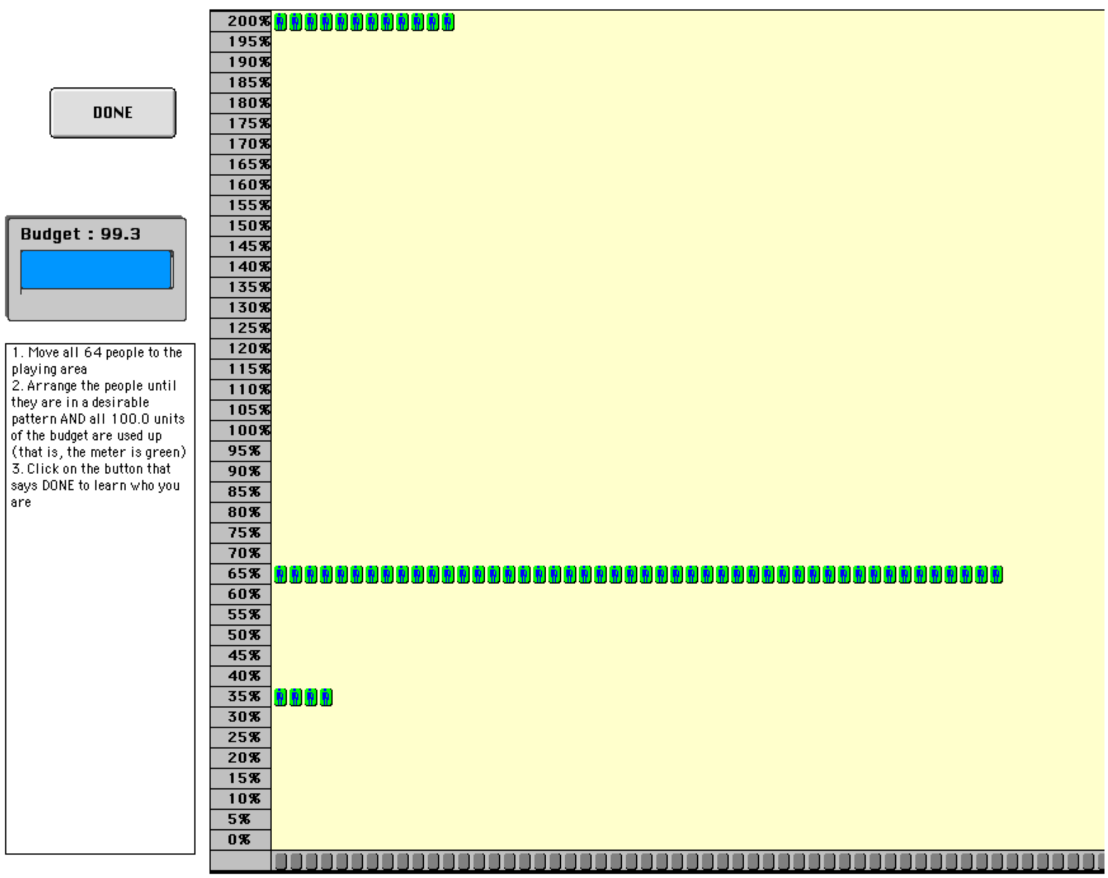

Using the mouse, the user can place the people in different rows, forming patterns against the vertical axis. Thinking of the number of people in a row divided by the total number of people as a probability, it can be seen that each pattern is equivalent to a probability distribution over levels of wealth. When the user begins interacting with the tool, all the people are in the reserve area and the budget meter (explained below) does not display a value. The user is told that she is represented by one of the people, but that all people look identical and there is no way to tell in advance which person she is. Given this information, the user is instructed to use the tool to create patterns that she would happily have apply to her own retirement income. The user can then place all the people on the playing field and arrange them into patterns against the income percentages on the vertical axis.

Each distribution that can be made with the Distribution Builder has an associated cost that is displayed on the budget meter. This cost is <u>not</u> the expected value of the probability distribution, but rather the amount of a hypothetical 100 unit budget that would be required to achieve that distribution of wealth using the cheapest possible dynamic investment strategy. When using the Distribution Builder, the user cannot select a final pattern that does not use 100 units of the budget.

In the application shown in Figure 1 the most conservative distribution that uses up the budget is achieved by placing all the people in the 65% row (which corresponds to investing all funds in a risk-free account). From this point, a little downside risk is rewarded with even greater upside possibilities. For instance Figure 1 shows a case in which (1) 4 people were moved from the risk-free 65% row to the 35% row and (2) 12 people were moved from the 65% row to the 200% row nonetheless leaving a small part of the budget unused.

For some purposes, the use of the tool ends when the user has selected his or her preferred feasible distribution. However, in some contexts it proves useful to include a second stage that simulates the realization of a specific outcome in order to help the user better understand the nature of probabilities. In the example shown in Figure 1, once the user decides on a desirable pattern, she can submit it to learn which of the people she is, and experience the process of learning how her retirement investment turned out. In this mode, after the user submits a distribution, the people begin to disappear from the board one by one until the only the only one left is the one representing the participant. This discrete representation of probability, in which the participant can envision herself as one of a number (here, 64) of people, should appeal to humans' preferential understanding of probabilities as frequencies[11].

2. Pricing a Probability Distribution

A key feature of the Distribution Builder is the pricing of probability distributions in a manner consistent with equilibrium in capital markets. We assume that the investor will choose combinations of broad asset classes and hence can achieve higher expected returns by taking on market-wide risk. To represent such trade-offs we utilize an Arrow-Debreu framework and procedures of the type developed by Dybvig[12]. A method for determing the underlying state prices is described in section 4. Here we focus on the use of such prices.

Consider an investor who is concerned only with the distribution of wealth at a specified horizon date H. We assume that her utility is a function solely of wealth at that date.

To simplify the analysis we assume that there are N mutually exclusive and exhaustive states of the world at the horizon date, and that each of the states has a probability of taking place equal to 1/N. The investor's *ex ante* measure of the desirability of a probability distribution is its expected utility, computed by weighting the utility of each possible outcome by its probability.

The investor has a given budget B and wishes to obtain a probability distribution of wealth that will maximize her expected utility without exceeding her budget.

We assume that there is a market in which one can obtain claims on wealth in the states and that the market is sufficiently complete that it is possible to arrange to obtain any given amount of wealth in one state and none in any other. The cost of obtaining $1 in state i is p_i. At least some prices are different, but we allow for cases in which two or more states have the same price. Without loss of generality, states will be numbered in order of increasing prices. Thus $p_i<=p_{i+1}$ for all i. The vector of these Arrow-Debreu state prices, $[p_1, p_2, ..., p_N]$ will be denoted **p**.

Consider an investor who desires a distribution of wealth D in which there is probability n_a/N of receiving wealth w_a, probability n_b/N of receiving wealth w_b, and so on, where n_a, n_b,... are integers. We may represent such a distribution by a vector of N wealth values in which n_a values are equal to w_a, n_b are equal to w_b and so on. For reasons that will become clear, we choose to arrange these values in order of decreasing wealth values. Thus $w_i>=w_{i+1}$ for all i. The vector of wealth values, $[w_1,w_2,...,w_N]$ associated with distribution D will be denoted **w**. Given our convention, there is a one-to-one mapping between the distribution D and the wealth vector **w** in the sense that for any distribution D there is a given wealth vector **w** and for any wealth vector **w** there is a given distribution D.

To obtain a set of payoffs with a given distribution D it is only necessary to assign each of the N wealth values in **w** to one of the N states of the world. We will call such an assignment an *investment strategy*. To determine the cost of any strategy one simply multiplies the price in each state times the wealth to be obtained in that state and sums the resulting products for all the states.

Clearly, there will be many possible ways to obtain a given distribution D and their costs may differ. We assume that the investor prefers to obtain a given distribution D using the strategy with the lowest cost. The goal is to find such a strategy and compute its cost. In this section we show how to compute the cost of such a strategy, in section 4 we discuss a procedure that can derive actual investment rules to achieve a desired strategy.

Consider an investment strategy in which w_i is assigned to state i, recalling that the states have been numbered in order of increasing prices and that the desired wealth values have been arranged in order of decreasing wealth. The cost of this strategy will be C = **p'w**. Importantly, there is no other investment strategy that will provide the distribution represented by **w** at a lower cost, although there may be others with the same cost.

To see why C = **p'w** is the lowest cost for which the distribution represented in **w** can be obtained, consider the conditions that would make a strategy not least-cost. Assume that for two states i and j, $p_i<p_j$ and $w_i<w_j$. The cost associated with obtaining w_i and w_j is $p_iw_i+p_jw_j$. But this can be reduced by switching the two wealth levels, so that w_j (the larger value) is obtained in state i (the cheaper state) and w_i(the smaller value) is obtained in state j (the more expensive state). Hence any strategy that allows for this kind of re-arrangement cannot be least-cost.

Now consider the manner in which the desired distribution was mapped onto states in our procedure. Since prices are non-decreasing in state number and wealth is non-increasing, there will be no cases in which any such re-arrangement can be used to lower total cost. Hence our procedure will always provide an investment strategy that is least cost.

Unless prices are strictly increasing in state number and wealth levels strictly decreasing, there may be alternative investment strategies that are least-cost, but under the assumption that utility is not state-dependent, the investor will be indifferent among all such strategies.

Clearly, a necessary condition for solving the investor's problem is the choice of a least-cost investment strategy. The design of the distribution builder restricts the investor's attention to such strategies. In this sense the investor is provided with investment expertise, allowing her to avoid strategies that are clearly suboptimal.

3. Inferring Attributes of a User's Utility Function

Implicitly, the user of the Distribution Builder is presented with a set of N Arrow-Debreu state prices and asked to choose a wealth for each one. If the goal is to determine an optimal investment strategy for the user this may be sufficient information. However, if other choices are to be made for the individual or if the information is to be used for calibrating models of equilibrium it is useful to interpret the resulting set of choices as consistent with the maximization of the expected value of a utility function of wealth and to infer some of the attributes of that function. Here, following Dybvig, we show how attributes of a user's utility function can be inferred from the distribution chosen.

Let u(w) represent the user's utility u as a function of wealth, w. The goal is to maximize the expected value of u(w) subject to the constraint that **p'w**=B.

Assume that the investor's utility function is smooth[13]. Let π_i be the probability of state i. To maximize u(w) subject to the budget constraint requires the satisfaction of the first order conditions that:

$$\pi_i u'(w_i) = kp_i \text{ for each state } i$$

where $u'(w_i)$ is the marginal utility[14] of w_i and k is a constant.

Since we have assumed that every state is equally probable, this can be written as:

$$p_i = \kappa u'(w_i) \text{ for each state i}$$

where $\kappa = (1/N)/k$.

Thus under the assumed conditions we may interpret p_i, the price of state i, as a constant times the user's marginal utility for the wealth w_i selected for that state. Thus the chosen distribution provides points on the investor's utility curve.

To illustrate, consider an investor with a power utility function (which exhibits constant relative risk aversion)[15]. Such a function has the form:

$$u(w) = w^{(1-g)}/(1-g)$$

giving a marginal utility of wealth of:

$$u'(w) = w^{-g}$$

Taking logarithms of both sides and writing the relationship for a specific state i gives:

$$ln(u'(w_i)) = -g\ ln(w_i)$$

As we have shown, the first order condition implies that:

$$p_i = \kappa\ u'(w_i)$$

which can be written as:

$$ln(p_i) = ln(\kappa) + ln(u'(w_i))$$

Combining the two equations gives:

$$ln(p_i) = ln(\kappa) - g\ ln(w_i)$$

We thus conclude that a user who maximizes the expected utility of wealth and who has a power utility function will select a distribution for which there is a linear relationship between *ln*(p_i) and *ln*(w_i), with the slope of the line equal to the negative of the exponent g in the underlying utility function.

There is, of course, no reason to believe that all users have power utility functions, nor that they will choose distributions that maximize the expected utility of this or any other utility function. Indeed, the relationship between **p** and **w** can be examined to assess the degree to which the user's choice conforms to maximization of *any* specified type of utility function. Nonetheless, given the prominence of the power utility function in the literature on lifetime consumption and investment planning, it seems especially relevant to investigate the extent to which the relationship between *ln*(**p**) and *ln*(**w**) for a user's choices can be approximated by a linear function.

Other questions can also be addressed by examining the relationship between **p** and **w**. For example, the maximization of a smooth utility function with a continuous first derivative requires that there be a one-to-one mapping between state prices and the associated levels of wealth. A kink in a user's utility function may be revealed by the choice of the same wealth in states with two or more different prices.

While the results that can be obtained with a Distribution Builder are limited in scope, they may well shed some light on investor preferences – light that is badly needed for both positive and normative applications.

4. Using a Binomial Process to Generate Prices and Determine Strategies

Thus far we have not indicated the way in which the set of state prices **p** utilized in an experiment might be chosen. In doing so, the goal is to utilize a set of prices that presents the user with trade-offs similar to those associated with actual investment markets. This section shows how a simple return-generating process can be used to generate a set of Arrow-Debreu prices and to also provide the rules for an investment strategy that can provide the chosen distribution at least cost.

4a. Characteristics of the Return-generating Process

A simple way to generate a set of Arrow-Debreu prices rests on the assumption that stock market returns follow a *binomial process*[16] in which there are two possible states of the world in each of a number of periods. We assume that the investor is allowed to allocate his or her assets between the stock market and a riskless security in each of H periods and that there are no transactions costs associated with any reallocation between these assets from period to period. To make the process even simpler, we assume that both the riskless rate of interest and the distribution of returns on the stock market are constant from period to period. In other words, we assume that returns are *independent and identically distributed* (I.I.D).

In any period, if the state of the world is that the market is *up*, \$1 invested in the riskless asset will grow to have a value of $\$v_r$, and \$1 invested in the stock market will grow to have a value of $\$v_u$. If the state of the world is that the market is *down*, \$1 invested in the riskless asset will still grow to have a value of $\$v_r$, but \$1 invested in the stock market will fall to a value of $\$v_d$. Finally, we assume that the two states of the world (up and down) are equally probable.

4b. Computing Arrow-Debreu Prices

Consider an investor who wishes to have \$1 at the end of a period if and only if the market is up. Assume that she may take either a long or a short position in one or both assets. Then, it will be possible to find a strategy using the two investments that will produce the desired payments. One only needs to solve the set of simultaneous equations:

$$x_r v_r + x_s v_u = 1$$
$$x_r v_r + x_s v_d = 0$$

where x_r represents the dollars invested in the riskless asset and x_s the dollars invested in the stock market. The cost of the resulting strategy (x_r+x_s) is the cost today of achieving a payment of \$1 at the end of one period if and only if the state of the world is up. We denote this p_u (the price of \$1 if the state is up). Replacing the right-hand side of the simultaneous equations with 0 and 1 provides the strategy that will provide \$1 if and only if the state is down. Its cost will be denoted p_d (the price of \$1 if the state is down).

To illustrate, assume that a period is one year and that all returns are in real terms. Let v_r=1.02, v_u=1.22 and v_d=0.92. Then p_u = 0.3268 and p_d = 0.6536. These are the state prices implicit in an economy in which the real rate of interest is 2%, the expected real return on the stock market is 5%, and the standard deviation of the real return on the stock market is 15% -- values not unlike those frequently used for projections made by academics and investment professionals[17].

Our simple one-period market is *complete* in the sense that the available securities allow the purchase of any set of outcomes over states at known prices. In such a market it is reasonable to assume that there are no possibilities for *arbitrage* (the ability to find an investment strategy that costs nothing to undertake, will provide positive income in one or more states and will provide negative income in no states). If the market is indeed arbitrage-free, each security will sell for an amount equal to the sum of the products of its payoff in each state times its state-price.

Now consider a multi-period setting in which there are H periods, each of which has the same distribution of outcomes. The simple two-branch tree process is now a substantial tree. There will be 2^Hdifferent paths through the tree and hence potentially different wealth levels for different investment strategies. Let N (=2^H) be the number of such paths. We seek p_i, the cost today of obtaining $1 at the horizon if and only if path i is realized.

In this case the price for a path can be determined directly once the number of "up" branches along the path has been specified. Let nu_i be the number of such branches along path i and nd_i (=H-nu_i) the number of down branches. If there are no multi-period arbitrage opportunities, the cost of receiving $ if and only if the path occurs will be:

$$p_i = p_u^{nui} p_d^{ndi}$$

since this will be cost of obtaining the payoff by using one-period investments as the path develops.

The relationship may be written in terms of the number of up branches on the path:

$$p_i = p_u^{nui} p_d^{(H-ndi)}$$

or as:

$$p_i = p_d^H (p_u / p_d)^{nui}$$

In our example, $p_u<p_d$, hence the parenthesized ratio is less than one, indicating that the price of a path will be smaller, the greater the number of up markets along the path. Note that all paths with the same number of up branches will have the same price. Thus, although there will be 2^H different paths, there will be only H+1 different prices (this follows from our assumption that the binomial process is I.I.D.). To maintain generality, we consider each path a different "state of the world" at the horizon, so that the number of states (N) will equal 2^H.

Recall that we wish to number states so that p_i increases with i. To do this requires only that we number the states in order of decreasing nu_i, so that $nu_i >= nu_{i+1}$ for all i. Thus state number 1 will represent a path in which the market goes up every period, states 2 through H+1 will represent paths in which the market goes up in one period and down in the other periods, and so on. Note that the assignment of numbers to paths is not unique. This implies that more than one investment strategy may provide a given distribution of terminal wealth for the same least cost amount.

4c. Finding a Dynamic Strategy

The Distribution Builder allows a user to construct any distribution. It is then priced using Arrow-Debreu prices. Eventually the user chooses a distribution which uses up her entire budget -- a distribution that is presumed to be preferred to all other feasible distributions. As indicated earlier, for some purposes this is all that is needed. In other cases it may be important to find an actual strategy that can provide the chosen distribution. Such a strategy will specify an initial mix of stocks and the riskless asset and a set of rules for changing this mix, depending on the path followed by the stock market up to each date until the horizon is reached.

Given a distribution D and a least cost implementation **w**, it is possible to determine the precise dynamic strategy that will produce the vector **w**. This can be done in one stage by solving a large linear programming problem[18]. Alternatively, the solution can be obtained by folding back the tree from its terminal nodes[19]. For each pair of terminal nodes with the same predecessor, the required amounts invested in the riskless asset and stocks can be found by solving the two simultaneous equations in two unknowns using the desired terminal wealth levels as the right-hand side. This provides the amount of money required at each of the nodes for period H-1. Once this has been done for all pairs of terminal nodes, the procedure can be repeated for the each of the pairs ending in period H-1, using the amounts determined in the prior step. The process is then repeated until the initial node is reached. The amount of money required at time 1 will, of course, equal **p'w**, which can be computed directly, as we have shown. However, the added information provides a complete set of instructions for allocation between cash and stocks at each node in the tree, and thus constitutes a detailed dynamic strategy.

The dynamic strategy required to obtain a chosen distribution may be simple or complex. In the current setting, in which returns are I.I.D., a *constant mix strategy* (with the same percentages invested in the two assets at all times and circumstances) will be optimal if an investor's preferences can be represented by a power utility function. As we have shown, such preferences will lead to the choice of a distribution for which there is a linear relationship between $ln(\mathbf{p})$ and $ln(\mathbf{w})$. Departures from such a relationship are especially interesting since they will generally imply a preference for outcomes that will require strategies that are truly dynamic, requiring changes in allocations of funds among major asset classes as the investor's wealth changes[20].

5. Summary and Suggestions for Further Research

We have described an approach that can be used in either experimental or practical applications. In either case, the goal is to obtain information about an individual's preferences based on his or her choice from alternative distributions of outcomes with the same cost. Key to the procedure is its focus on realistic alternatives that reflect the manner in which capital markets can evolve.

Our goal has been to illustrate the approach with a relatively simple example in the hope that others will apply and extend it. To conclude we briefly outline some possible avenues for further research followed by a few caveats.

Our example employed 64 people in order to utilize a binomial model of asset price behavior. One could of course use a larger number of periods in the return-generating process, expanding the number of people to 128, 256 or any power of two, thereby providing the user with more degrees of freedom. Alternatively, one could employ a discrete approximation to a continuous distribution of Arrow-Debreu prices[21]. Among other things, the latter approach would allow for a presentation involving 100 people, which would have the advantage of equating the number of people associated with an outcome with the probability of that outcome.

While it may prove convenient to assume that stock prices are independent and identically distributed (I.I.D.), there is no need to do so. For example, prices could be assumed to follow a binomial process in which the expected stock return, variance of stock return and/or riskless rate of interest could be dependent on prior events. More complex stochastic processes[22] could also be utilized, as long as sufficient instruments exist to span the set of outcomes so that Arrow-Debreu prices can be calculated.

While the approach is rich in possibilities, it is not without limitations. Our example required the user to focus on a single outcome (income in retirement). One can envision variations that would utilize two outcomes (for example, savings per year prior to retirement and wealth at retirement)[23] but extensions designed to estimate characteristics of a user's full multi-period utility function would require restrictions on the assumed nature of user preferences.

Finally, there are serious questions about the ability of a user to fully understand the trade-offs presented by the Distribution Builder. In some cases understanding might be better if the underlying Arrow-Debreu prices were presented explicitly. In other cases, a user might need to engage in several experiments before fully understanding the nature of the available alternatives. Hopefully, behaviorists will be able to perform experiments that can lead to presentations that more efficiently obtain information about users' true preferences.

Footnotes

* William F. Sharpe is STANCO 25 Professor of Finance, Emeritus at Stanford University and Chairman, Financial Engines, Inc.. Daniel G. Goldstein is Director, Products Division at the Fatwire Corporation. Much of this work was performed when Goldstein and Blythe were associated with the Center for Adaptive Behavior and Cognition at the Max Planck Institute in Berlin. An early version of this paper was presented at the Third International Stockholm Seminar on Risk Behaviour and Risk Management in June, 1999.

1. Sharpe, 1964

2. Markowitz, 1952

3. See, for example, Hong and Wang, 2000

4. See, for example, Epstein and Zin, 1989 and Constantinides, 1990

5. See, for example, Barberis, 2000

6. See, for example, Hong and Wang, 2000

7. For an overview, see Campbell, 2000

8. Kahneman and Tversky, 1979

9. For a discussion of this and other aspects of interest in this context, see Shefrin, 1999

10. For a formal model, see Brennan and Solanki, 1981

11. For a discussion, see Gigerenzer, 1994.

12. The basic concept of state-prices was developed in Arrow, 1964 and Debreu, 1959. The computation of the least-cost of a distribution is presented in Dybvig, 1988.

13. More precisely, that the first derivative is continuous.

14. that is, the derivative of u(w) with respect to w at w=wi.

15. For a discussion of the properties of this and other forms of utility functions, see Huang and Litzenberger, 1988

16. initially developed in Sharpe, 1978 and greatly expanded in Cox, Ross and Rubenstein, 1979

17. See Welch 1999

18. See, for example, Sharpe, 2000

19. See, for example, Sharpe, Alexander and Bailey, 1999, pp. 616-622.

20. For a related discussion, see Perold and Sharpe, 1988

21. For example, in the limit (as the number of periods H goes to infinity), the Arrow-Debreu prices for an IID process converge to a lognormal distribution.

22. such as trinomial processes in which there are three possible states of the world in each period.

23. For example, the user could be presented with a grid in which one outcome is plotted on the horizontal axis and the other on the vertical axis. The user could then place people on cells in the grid.

References

Arrow, Kenneth J. "The Role of Securities in the Optimal Allocation of Risk-bearing," Review of Economic Studies, April 1964, pp. 91-96

Barberis, Nicholas, "Investing for the Long Run when Returns are Predictable," Journal of Finance, Vol. 55, no. 1 (February 2000), 225-264.

Brennan, M.J. and R. Solanki, "Optimal Portfolio Insurance," Journal of Financial and Quantitative Analysis, 16, no. 3 (September 1981): 279-300.

Campbell, John Y., "Asset Pricing at the Millennium," Journal of Finance, Vol. 55, no. 4 (2000) 1515-1568.

Constantinides, George, "Habit Formation: A Resolution of the Equity Premium Puzzle," Journal of Political Economy 94, (1990) 842-862.

Cox, John C., Stephen A. Ross and Mark Rubinstein, "Option Pricing: A Simplified Approach," Journal of Financial Economics, 7, no. 3 (September 1979), 229-263.

Debreu, Gerard, Theory of Value, The Cowles Foundation Monograph 17, 1959

Dybvig, Philip H., "Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market," *The Review of Financial Studies*, (1988), Volume 1, Number 1, 67-88.

Epstein, Lawrence and Stanley Zin, "Substitution, Risk Aversion and the Temporal Behaviour of Consumption and Asset Returns, A Theoretical Framework," Econometrica 57, (1989) 937-969.

Gigerenzer, Gerd, 1994, Why the distinction between single-event probabilities and frequencies is relevant to psychology (and vice versa). In George Wright and Peter Ayton, Eds., Subjective probability(Wiley, New York).

Hong, Harrison and Jiang Wang, "Trading and Returns under Periodic Market Closures," The Journal of Finance, Vol. 55, no. 1 (February 2000), 297-354.

Huang, Chi-fu and Robert Litzenberger, Foundations for Financial Economics, Prentice-Hall (1988).

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Choices Involving Risk," Econometrica, 47 (1979), 263-291.

Markowitz, Harry, "Portfolio Selection," Journal of Finance 7 (1952) 77-91.

Perold, Andre F., and William F. Sharpe, 1988, "Dynamic Strategies for Asset Allocation", Financial Analysts' Journal, Jan/Feb, 16-27.

Sharpe, William F., Investments, 1st Edition, Prentice-Hall, 1978.

Sharpe, William F., "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," Journal of Finance, 19 (1964) 425-444.

Sharpe, William F., Gordon J. Alexander and Jeffery V. Bailey, Investments, 6th Edition, Prentice-Hall, 1999.

Sharpe, William F., Macro-Investment Analysis: Prices, Dynamic Strategies, 2000 www.wsharpe.com/mia/prc/mia_prc3.htm#dynamic

Shefrin, Hersh, Beyond Greed and Fear: Understanding Behavioral Finance and the Psychology of Investing, 1999.

Welch, Ivo, "Views of Financial Economists On The Equity Premium And Other Issues," 1999, available at: http://welch.som.yale.edu/academics/#mktsurvey