Using Asset Allocation to Protect Spending

Philip H. Dybvig*
Washington University in Saint Louis

ABSTRACT

Management of portfolios from which funds are withdrawn involves two strategic decisions at the fund level: asset allocation and the spending rule. Asset allocation and the spending rule are traditionally linked, for example, to preserve capital on average, but there should be a closer dynamic linkage between the two. This article describes a new protective strategy that links spending and asset allocation in a way that preserves spending power when the market is down while still participating significantly when the market is up. This strategy is similar to Constant Proportions Portfolio Insurance in that part of the fund is kept in safe assets to preserve the value needed for continued expenditures. Like portfolio insurance, the strategy outperforms traditional strategies when markets are persistently up or persistently down, but underperforms when whipsawed by repeated up-and-down moves (This is the cost of the downside protection).

This draft: October 30, 1998

*I am grateful for discussions with Dick Anderson of Kennedy Capital, Bill Marshall of NISA, André Perold of Harvard, and Todd Petzel of Commonfund, and I am grateful for support from Commonfund. This article was written in connection with the Commonfund Prize and is intended to communicate to practitioners the research honored by the Prize. Commonfund (previously named The Common Fund) is a nonprofit organization that is a cooperative undertaking of educational institutions to improve the productivity of endowments. The Commonfund Prize was awarded in celebration of the organization's 25th anniversary.

Introduction

Management of educational endowments involves two strategic decisions at the fund level: asset allocation and the spending rule. Asset allocation is the decision of how much of the endowment should be placed in various broad asset classes, such as domestic equities, foreign equities, and government bonds. The spending rule is the decision of how much should be removed from the endowment for current spending, and how much should be left in to fund future spending. Asset allocation and the spending rule are traditionally linked, but there should be a dynamic linkage between the two that is not traditionally made. The purpose of this article is to provide guidance on this overall strategy to the university officers responsible for asset allocation and the spending rule. The strategy proposed in this paper does a better job of protecting capital than current practice, but still participates significantly in up markets. This strategy is based on recent breakthroughs in the theoretical literature and specifically on Dybvig [1995].

1 Current Practice and How To Do Even Better

Conceptually, endowment management concerns an alternating sequence of spending and asset allocation decisions. At the beginning of a year (or other planning period), we look at the current size of the endowment and we decide how much to extract for expenditure this period, and how much to leave in the endowment to invest for future spending. This is the spending problem. Then, we take what remains in the endowment and invest in some mix of different asset classes. This is the asset allocation problem. After a year's realized return on investments, we look at what is now in the portfolio, including the endowment at the beginning of the year adjusted up or down based on the whether fortune smiled on our investments (depending on our skill and our luck), plus other sources or sinks of income such as net new donations, money budgeted but not spent, tuition beyond what was expected, etc. Then we return once again to the asset allocation problem of deciding

how to allocate our investments to different asset classes.

In recent decades, universities have been increasingly sophisticated in linking budgets and investment policy. There is necessarily a linkage between the two, since what funds are available to one is given by what was left over from the other. Having more of a link between the two once had a bad name because of the archaic limit on spending more than dividends plus interest (see Keane [1996] and Murray [1996] for more about the history). Current practice tends to link budget proportions to average long-term returns to maintain principal on average. Nonetheless, this link is static and does not have dynamic feedback between the two like the link I am proposing. There is also a serious problem with the usual calculation that is done for preservation of capital, but that is a topic for another paper.[1]

To see an example of the performance of a strategy that is a first approximation to current practice, consider Figures 1 and 2. They backtest the performance of a strategy of putting half of the portfolio in large-cap stocks and half in long government bonds with a spending rule of 4.5% per year.

A slightly more complex strategy is a better representative of current practice. Asset allocation may come from proportions dictated (within a band) by an asset pricing model such as the CAPM or APT, which helps us to quantify the trade-off between risk and return. The spending rule is smoothed compared to the historical rule (for example by basing spending on a 5-year moving average of endowment values or unit values instead of the current value), and the overall level of expenditure is set using a rule-of-thumb intended to preserve principal on average.

To see the effect of smoothing on spending, consider Figures 3 and 4. They

[1]Saying that capital is preserved on average is not really good enough. For example, betting the entire endowment on a horse race with fair odds preserves capital on average, but does not conform to any reasonable notion of preservation of capital. If the mean return of the portfolio is 10% with a standard deviation of 30%, expected inflation is 3%, and spending is 4.5%, the usual calculation would say capital is preserved, since spending is less than the mean return net of expected inflation. However, the real value of the endowment tends to zero almost surely over time, since the value looking forward to a date far in the future is similar to the payoff in a lottery ticket. For details, see Dybvig and Heston [in preparation].

Figure 1: Performance of a simple portfolio strategy and spending rule: spending (smoother curve) and inflation-adjusted portfolio value

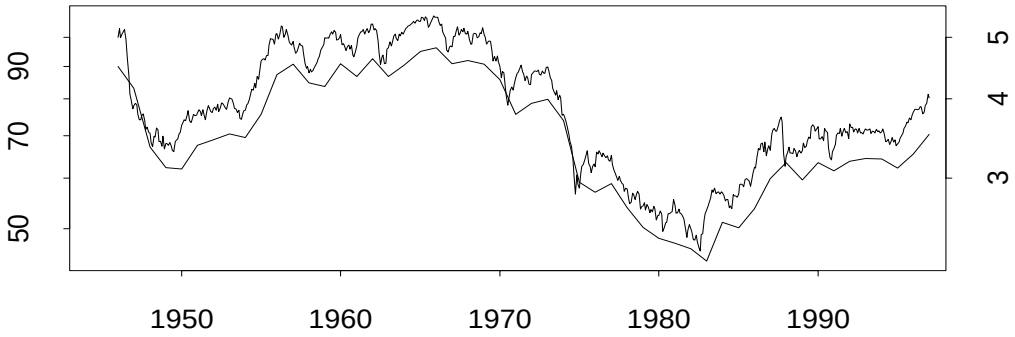

This shows the inflation-adjusted performance of a traditional policy having half of the portfolio invested in large-cap stocks and half in long government bonds, with expenditure of 4.5% per year, for the post-war period. The graph shows wealth (the size of the endowment—left scale) and spending (right scale). Spending is the smoother curve only because it is changing only once a year.

Figure 2: Corresponding portfolio composition: bonds (lower curve) and wealth = bonds + stocks (upper curve).

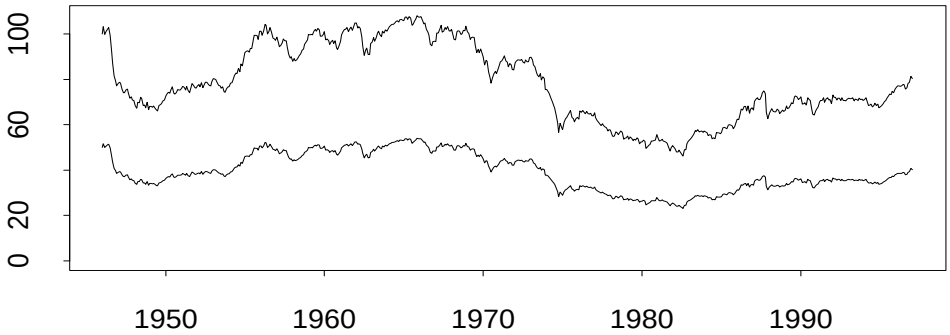

This is the corresponding portfolio strategy. Since most of the uncertainty is in the value of the stock portion of the portfolio, funds are moved from the stock portfolio to the bond portfolio as the market goes up, and from the bond portfolio to the stock portfolio as the market goes down. In this way, the proportions in the two asset classes remain constant.

are based on the same spending rule as in Figures 1 and 2 except that spending is based on a 5-year aveage of the endowment rather than the current endowment. While this does reduce the typical magnitude of year-to-year changes, it does not change the overall shape of the expenditures over time. In principle, there is even a possibility under this rule of a much larger dip than even the traditional rule if we should see a more persistent decline in the stock market than has occurred in our post-war sample. This is because expenditure following a persistent decline could be at much greater than the 4.5% rate, resulting in a significant erosion of capital.

While the smoothed spending rule may reasonably serve the needs of an institution in many market environments, it nevertheless is *ad hoc* with little basis in theory. This practice of smoothing spending can be improved by understanding that asset allocation and the spending rule are linked by the nature of the underlying expenditures we want to protect. I specifically wish to consider strategies that can ensure that spending will never decline, or will never fall faster than a pre-specified rate of decline. Figures 5 and 6 provide backtests of the strategy I am proposing.

Unlike current practice, the new strategies avoid significant spending cuts while still participating to a significant extent in rising markets. The large dip in real spending that occurred in the mid-80s for the traditional strategy would not have occurred in the strategy I am proposing. That is because the portfolio strategy would have protected the ability to spend at the existing level by moving out of risky assets and into safer investments as the market fell. This feature allows us to avoid damage done by falling budgets while still participating significantly in up markets. We will discuss what is happening in Figures 5 and 6 once we have described the strategy more fully.

One important feature of the proposed strategy is that spending rates are less than what is now common. This is necessary because it is not possible to maintain for certain a perpetual spending level that is higher than the current endowment times the riskless rate. A modification of the strategy (described in detail in Section 4) allows us a higher initial spending rate but the cost is a declining process. There is no free lunch; the higher the initial spending rate, the faster the spending rate must be able to decline, and the more the strategy looks just like the traditional strategy without protection,

Figure 3: Performance of a portfolio strategy and spending rule with smoothing: spending (smoother curve) and real (inflation-adjusted) portfolio value

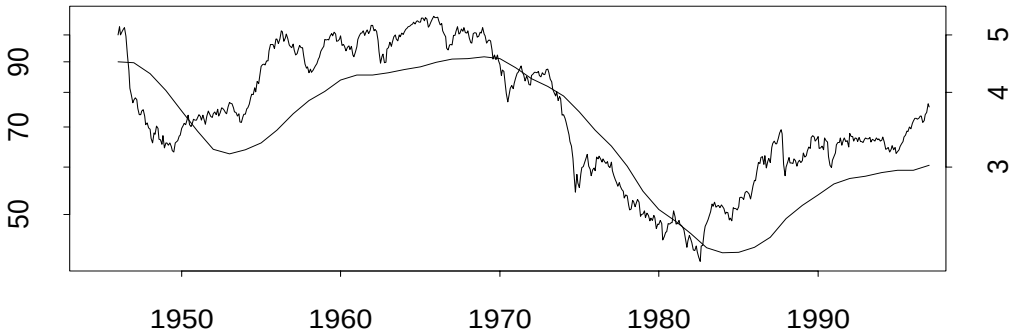

This shows the constant-dollar performance of a more modern policy over the post-war period. As before, half of the portfolio is invested in large-cap stocks and half in long government bonds. Expenditure, however, is given by 4.5% per year of a 5-year weighted average of the portfolio value. The graph shows the wealth (left scale) and spending (right scale). Spending is even smoother than in Figures 1 because it is based on a 5-year moving average of the portfolio value instead of the portfolio value at a point in time. Nonetheless, the overall pattern of spending over time is not much different from what it is in the traditional rule.

Figure 4: Corresponding portfolio composition: bonds (lower curve) and wealth = bonds + stocks (upper curve).

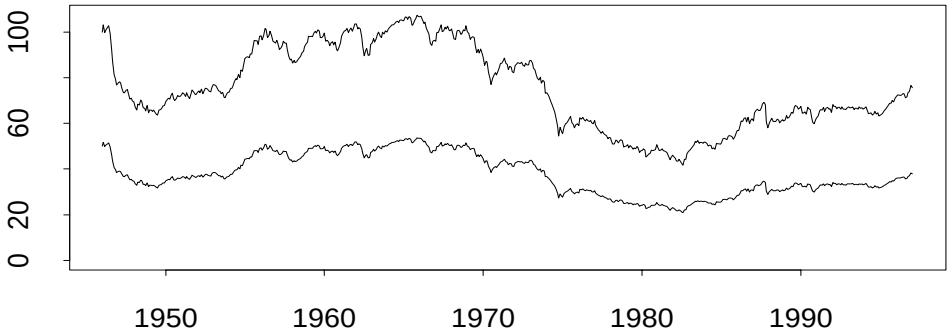

As in Figure 2, this strategy maintains equal proportions of the portfolio's value in the two asset classes.

Figure 5: Performance of the proposed expenditure-protected policy: spending (smoother curve) and real portfolio value

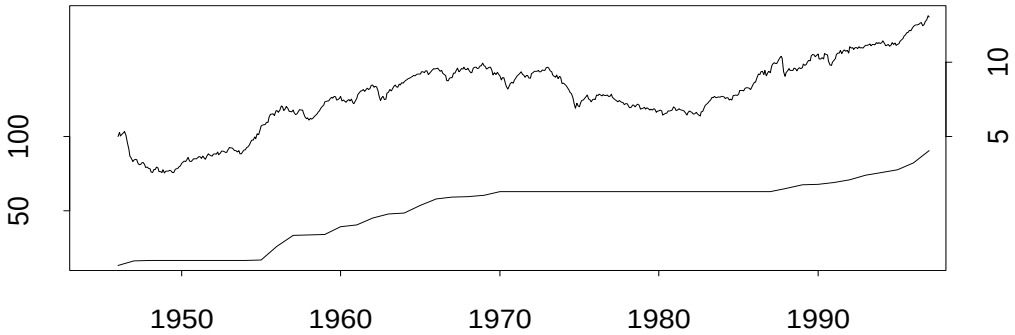

This shows the constant-dollar performance of an expenditure-protecting policy, for the same post-war periods examined in Figures 1 and 3. For a real riskless investment, we use short Treasuries, which do a good job of tracking inflation over time. This strategy avoids times of particular duress in which real spending falls.

Figure 6: Corresponding portfolio composition: bonds (lower curve) and wealth = bonds + stocks (upper curve).

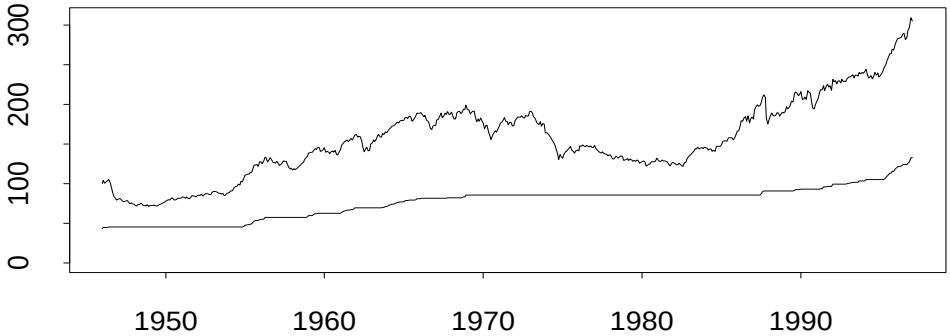

Unlike in the traditional strategies, the funds earmarked for protected expenditures stay in the less risky bonds as the market falls, even if this means that bonds will represent a large proportion of the portfolio. As the stock market rises, on the other hand, wealth is moved from risky assets to the bond portfolio as needed to ensure that the overall portfolio does not become too risky. This transfer of funds to the bond portfolio corresponds to an increase in committed spending.

as illustrated in Figures 7 and 8. These two cases are extremes; in general, the higher the initial spending level, the smaller the guarantee we can afford.

2 Protected Expenditure

Central to our approach to coordinating asset allocation and the spending rule is the notion of protected (or committed) expenditure. At any point in time, we will have a minimal amount of future spending we plan to obtain from the endowment, no matter how badly the stock market performs. This is possible because, as the value of the endowment falls, we do not transfer more and more into risky investments as we would in a fixed-proportions strategy, and instead we keep enough in safe investments that will deliver enough cash flow to cover our future protected expenditure. To think of this another way, our portfolio is partitioned into a protected part and a risky part. The protected part contains riskless investments whose payoffs are designed to produce the planned protected expenditure over time, while the risky part is poised to capture the average expected returns above the riskless rate, which represent the compensation for taking on risk.

Protected expenditure could refer to actual commitments we have made (for example for salary or funding a program), or it could just represent our best judgment of the income pattern below which unacceptably large damage would be done. If we consider a fixed date in the future, then as the date approaches our protected expenditure at that date will usually rise, as we commit more and more to the expenditure at that date. The full commitment is normally made at the time the allocation moves from the arena of investments into the budgeting process.

Of course, it would be nice to be able to have protected expenditure at a very high level for all future dates. Unfortunately, we are subject to the constraint that the present value of our protected expenditure must be no greater than the value of our endowment. If we manage a portfolio worth \$100 million and the interest rate is 5%, then the highest constant commitment to spending we can make in perpetuity is \$5 million per year, and if we make

Figure 7: Performance of a diluted version of the proposed expenditure-protected policy: spending (smoother curve) and real portfolio value

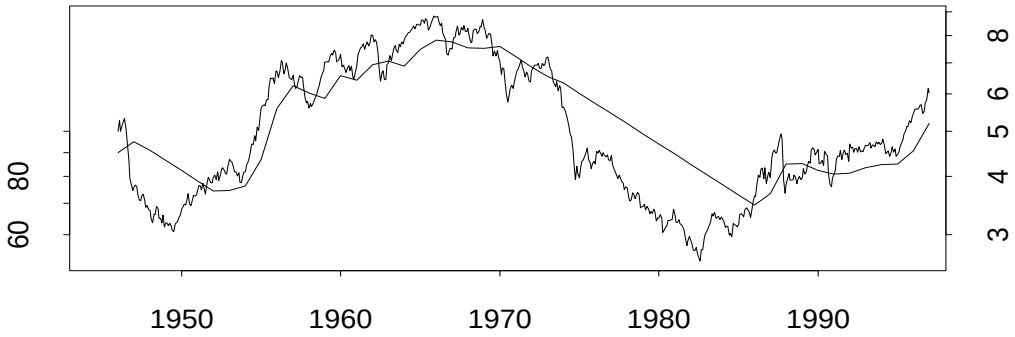

This shows the constant-dollar performance of an expenditure-protecting policy that limits the decline to 5%/year. This allows a much higher initial spending rate than in the fully-protected strategy in Figures 5 and 6, but the protection is not worth much and in fact the overall pattern of spending is uncannily close to the pattern in the traditional strategy in Figures 3 and 4.

Figure 8: Corresponding portfolio composition: bonds (lower curve) and wealth = bonds + stocks (upper curve).

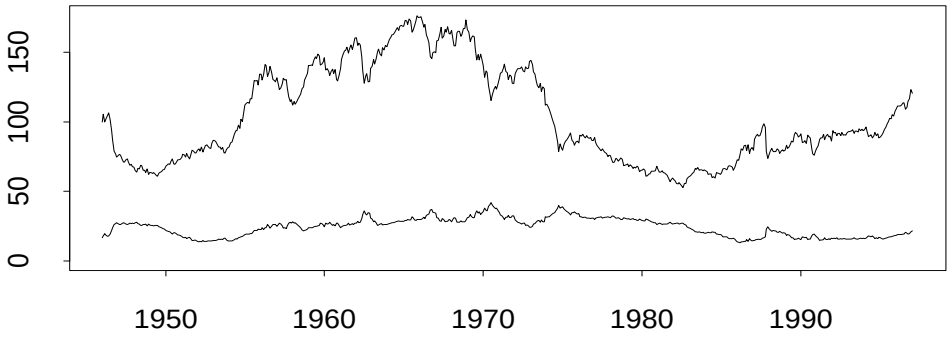

Given that consumption is allowed to fall significantly, the portfolio choice adjusts more similarly to Figure 2 than to the protected strategy in Figure 6, since it does not require much wealth to support the declining commitment.

that commitment we will be locked out of the stock market and the higher average returns available there. And, if we want protected expenditure to grow over time, an even lower level of initial expenditure would be available. In particular, it makes sense to think of expenditure that is protected in real (inflation-protected) terms, so the purchasing power stays constant. In this case, we can probably treat short Treasuries as a real (inflation-protected) riskless asset and expect a real return of say 3%. Given the new issues of indexed Treasury Bonds, it should be possible to lock in spending power with even more confidence than in the past. Nonetheless, Treasury Bills have historically yielded only a slight premium above inflation, and I am unsure just how high a real return we can rely on being available into the future.

Given that stocks tend to outperform Treasuries in the long term, why can't we spend in perpetuity a larger amount than $5 million using investment in stocks? The answer is that we can make a larger payout with some probability, *but there is also a significant probability of a shortfall.* Asserting otherwise says that we are certain that stocks will go up and that going long in stocks and short in the riskless asset is in effect a riskless arbitrage. This sort of cheerful optimism is a very appealing personality trait, but not a healthy attitude for an investment manager. This is exactly the sort of perspective that some people had one month before jumping out of the window during the market crash in 1929. Should we worry about something that happens so infrequently? Of course we should, since we are in the business of managing the endowment of an organization that is supposed to survive for a long time. It is tempting to say that we should not be worried about something that did not happen in the past 50 years, but 50 years is not such a large time to see all that can happen in a financial market, and we should be prepared to deal with whatever reasonable contingencies may occur. There is an additional reason to distrust drawing strong conclusions from the historical performance of the US stock market, which is that it has performed beyond expectations. Stock markets that were equally promising (or more promising) 100 years ago have performed less well or have disappeared entirely. We study the US stock market precisely because it is atypical in its survival and prosperity.[2] For example, Jorion and Goetzman [forthcoming]

[2]This idea of survivorship bias in measuring U.S. market returns was first introduced to me by Pete Kyle in the early 80's.

consider equity returns from a variety of countries from 1921 to 1996. The average U.S. real rate of return in their sample is 4.3%, while the average over all countries is only 0.8%. If we take these numbers seriously, it is hard to make the case that current endowment investment policies can be relied upon to preserve capital.

In general, there are a number of different ways we can set up the planning process for commitment of spending. We can commit a fixed dollar amount each year in perpetuity (as in the example above). We can adjust the amount for inflation, promising a larger dollar amount in each year, or we can make the commitment decline over time with the understanding that decreasing expenditure is not so devastating if it is gradual. Or, we may plan to commit spending for a fixed number of years in the future, extending our commitment each year only to the extent that is prudent given future circumstances. We will consider in more detail the different types of commitment and their implications in a later section. For now, we turn to the commitment-based model of asset allocation and the spending rule.

3 The Proposed Strategy

In Dybvig [1995], the strategy proposed here was derived as the solution to an optimization problem (given below as Problem 1). In fact, that analysis started with a description of the economic situation, which was formalized in a choice problem, and subsequently solved. For practice, however, the strategy itself is most important, so the strategy is our starting point and our main focus.

The proposed strategy is given in Table 1. At a point in time t, we have wealth (or size of the endowment) w_t, of which α_t is invested in the risky portfolio (say domestic equities or a favored mixture of various risky asset classes) and $w_t - \alpha_t$ is invested in the riskless assets (short Treasuries or better yet inflation-protected Treasuries). Spending is at an annualized rate c_t. The constant r is the riskfree rate (presumably the excess over inflation). The constants K and r^* are parameters of the strategy. They are chosen

Table 1: The Proposed Strategy: No Decline

Portfolio choice: The part of the endowment placed in the risky portfolio is

$$\alpha_t = K(w_t - s_t/r),$$

and the amount placed in riskless assets is the remainder

$$w_t - \alpha_t = s_t/r + (1 - K)(w_t - s_t/r).$$

Spending rule: Spending is normally kept constant and is increased only as is necessary to maintain

$$s_t \geq r^* w_t,$$

with $r^* < r + d$, where r is the riskless rate of interest and d is the permitted rate of decline That is, there is a critical proportion r^* of wealth below which spending can never fall. The plots in Figures 5 and 6 take $K = 1$ and $r^* = 1.5\%$/year.

at the outset. My preference is to choose the parameters by experimenting with simulated returns to obtain the most appealing range of results in a number of simulated draws. An alternative is to base the parameters on the theoretical model (Problem 1 discussed below) that gives formal justification of the strategy. Unfortunately, the theoretical model relies on preference parameters that must be supplied. I think it is easier to explore preferences by looking at simulated outcomes than by deciding on the parameters determining the mathematical specification of time and risk preferences. Point to my Web site *http://dybfin.wustl.edu/java/ratchet* to simulate the strategy; details are in Section 4.

In the strategy, we interpret s_t/r as the wealth required to maintain current spending forever, since this is the amount we would have to invest at the riskless rate r to generate interest of s_t each year. Therefore, the portfolio choice has the cost of maintaining current spending forever, s_t/r, invested in the riskless portfolio, and the remainder, $w_t - s_t/r$, invested in fixed proportions, K in the risky portfolio and $1 - K$ in the riskless portfolio. In the examples of Figures 5 and 6, we took $K = 1$ so the entire cushion is invested in the risky asset. Even taking $K > 1$ (more than 100% of the cushion in equities) need not ever lever the portfolio as a whole, provided $K \leq r/(r - r^*)$.

For the spending rule, we know we never decrease spending from the previous year, since that is a constraint of the problem. If fortune has smiled on us and the portfolio has increased in value over the year (or we received an unanticipated capital gift), then we increase spending only to the extent needed to maintain a spending rate of r^*. Over time, spending can vary between a minimum of this critical rate r^* and a maximum of r, depending on our luck and skill in our risky investments. If $s_t/w_t = r$, then the portfolio is fully invested in bonds. When $s_t/w_t < r^*$, you raise spending immediately to the level s'_t where $s'_t/w_t = r^*$. This new value s'_t is the new level of spending that can never decrease.

The investment policy is tied to the spending rule. Although implementation of the strategy does not require separate accounts, the rule works as if there are two separate accounts, one nondiscretionary account (the "committed amount") for funding the committed expenditure of current spending continued in the indefinite future, and another discretionary account (the "cushion") invested in fixed proportions in stocks and bonds. The spending rule keeps expenditure at a fixed level until we need to raise it to maintain a minimum spending rate. Therefore, when the portfolio value goes down, we maintain spending, but we increase spending when the portfolio value goes up enough. As the portfolio value goes down, a smaller and smaller proportion of the portfolio is at risk in the discretionary part, permitting us to maintain enough value in the committed part of the portfolio to keep from ever having to reduce spending.

Recall that Figures 5 and 6 show the performance (inflation-adjusted) of this strategy during the same post-war period examined in Figures 3 and 4. The spending under this policy is qualitatively quite different from the spending in the current practice, since it never falls. Also, the portfolio choice is quite different: the dollar amount left in riskless assets is much more stable, especially in response to falling equity markets. The late 70s and early 80s were a difficult time for many universities, and we can see one reason why from Figures 3 and 4. In the traditional strategy, a falling market implied falling real spending. By contrast, the strategy I am proposing would preserve spending. An important part of this strategy is retention of enough safe investments to ensure that maintenance of spending is always possible.

Table 2: Some properties of the solution to Problem 1

- Investment policy
  - Conceptually, separate the portfolio into committed and discretionary accounts.
  - Immunize the committed part to fund committed spending.
  - The discretionary part is invested in different assets in fixed proportions.
  - As committed spending increases, the funds needed to ensure committed spending in the future are transferred from the committed account into the discretionary account.
  - This investment policy is generally consistent with the CPPI policy analyzed by Black and Perold [1992].
- Spending rule
  - Increase future commitments in response to superior performance of investments.
  - Plan to preserve capital to an extent dictated by time and risk preferences.

Some qualitative properties of the solution are summarized in Table 2.

Some of the important distinctions between the traditional strategy and my proposal are contained in Table 3.

The Formal Choice Problem

The formal choice problem is given in Table 4. The exact meaning of all the symbols is probably unclear to anyone without formal training in probability theory and continuous-time finance. Nonetheless, it is still possible to get an intuitive feel for what is going on without that background, and some of the specifics are given in the caption to the table. The formal choice problem

Table 3: Comparison of the new strategy with the traditional strategy

- The new strategy has some qualitative properties that are similar to portfolio insurance.
- There is a much better worse-case scenario (for a persistently declining market) than the traditional strategy. In particular, spending will never decrease.
- There is a better best-case scenario since there is more participation in a persistently rising market.
- The new strategy is worse in up-and-down markets. This is the whip-saw effect that is familiar to users of portfolio insurance.

is almost the same as the famous problem studied by Merton [1973], except for the additional constraint that spending can never fall. The Appendix gives formulas spelling out the connection between the constants K and r^* in the formal solution and the parameters of the formal choice problem. The derivation of the solution is given in Dybvig [1995].

4 Getting Started

Before implementing the strategy I am proposing there are several essential steps:

1. Plan to follow the strategy over a period of years. In general, switching among reasonable strategies destroys value. See Dybvig [1988].
2. Choose the parameters of the strategy. In the simplest form of the strategy, this means picking the critical minimal spending rate r^* and the risky portfolio proportion K for the discretionary part. Two potential approaches are described below. The first approach uses simulations of the portfolio value to help us to make the trade-offs. This is the approach I am recommending, and there are instructions below on how

Table 4: A formal decision problem

Problem 1: Choose the *spending rule* s_t and *risky portfolio investment* α_t to maximize the *objective function* $E[\int_0^\infty u(s_t)\exp(-\delta t)dt]$ subject to

$$(\forall t' > t)s_{t'} \geq s_t \quad \text{(nondecreasing spending)}$$

and

$$(\forall t)w_t \geq 0, \quad \text{(nonnegative wealth)}$$

where w_t solves the *budget equation*

$$dw = w_t r_t dt + \alpha_t(\mu dt + \sigma dZ_t - r dt) - s_t dt$$

subject to the *initial wealth constraint* $w_0 = W_0$. The *utility function* is the power function $u(s_t) = s_t^{1-R}/(1-R)$.

The variable t indexes time and the choice variables are the spending rule s_t, which is the rate per unit time at which we are spending out of the endowment, and the risky portfolio investment α_t, which is the proportion of the endowment invested in risky assets. Both of these can depend on time and on the past realizations of portfolio returns. The objective function $E[\int_0^\infty u(s_t)\exp(-\delta t)dt]$ says that we want to maximize the sum ($\int$ indicates integral, the continuous analog of a sum) over time of instantaneous benefits $u(s_t) = s_t^{1-R}/(1-R)$, weighting benefits at t by a weight $\exp(-\delta t)$ that is declining over time. The parameter δ is the pure rate of time discount, and the larger its value the less importance we place on later spending. The first two constraints say that spending can never decline. The third constraint ensures we cannot borrow forever without repayment. Finally, the budget equation says how wealth changes. If invested just in the riskless investments we get the riskless rate r per unit time. For the amount α invested in risky assets, we get a mean return μ per unit time and random return σdZ_t, and we lose out on the riskless return r per unit time. Also, we take from the endowment whatever we spend at the rate s_t per unit time.

to perform the simulation by pointing your Web browser at a program I have written. The second approach takes the model more seriously as a literal description of what we are about, and the parameter choice is based on our preferences. I am not recommending this approach, but it is discussed below for those who would like to try it.

3. Develop a transition plan. It is likely that spending under the strategy I am describing will be less than under current rules (although at present we may have a good opportunity due to value increases in the bull market). Therefore, it is probably a good idea to phase in the new plan, for example, by moving a fixed proportion each year into the new program.

In the rest of this section, we describe the two methods for deciding on parameter values. The following section considers some other implementation issues.

Parameter Values from Simulation

Simulating investment results from the strategy is perhaps the best way to internalize the implications of the strategy and decide on how to choose the parameters of the strategy. The advantage of using simulation to evaluate parameter choices is that it does not require us to buy into the whole theoretical model or determine the preference parameters in the model. Any theoretical model is necessarily an abstraction of reality; this seems especially true of models of preferences.

To compare traditional asset allocation and spending rules with my proposed model, I suggest generating plots of the annual spending from endowment and portfolio value like those in this paper only using simulated portfolio returns and not just the historical returns as shown in this paper. This gives a more accurate picture of the range of performace scenarios that might reasonably happen in practice than looking at the single scenario in the historical record. I have prepared a program designed to perform the simulations I am describing; at the time of writing of this paper it may be found on the Web

at the URL *http://dybfin.wustl.edu/java/ratchet*. For any such simulation, I recommend using the same stochastic model of portfolio returns as in Problem 1 (Table 4), although it is also possible of course to include additional realistic properties of returns (such as stochastic volatility or time-varying risk premia).

Parameter Values from Underlying Preferences

One way to decide on parameters of the model is to take seriously the derivation of the target portfolio proportion K and the minimum spending rate r^* from the theoretical model. The advantage of this approach is theoretical precision.

The preference parameters we need to specify are the pure rate of time discount δ and the degree of relative risk aversion R that enter into the objective function in Table 4. The pure rate of time discount δ enters directly and the risk aversion R enters through the felicity function u. We can get some handle on these parameters by asking questions such as "What riskless percentage increase in spending would be just as attractive as a $50-50$ chance of no increase or a 10% increase?" This would be compared to the theoretical value x solving the equation $u((1+x)s) = .5u(s)+.5u(1.1s)$ (which, for the u functions we are considering, does not depend on w but does depend on R). Once we know R, we can determine δ by asking a question such as "What percentage increase in spending both this year and next would be just as attractive as a 10% increase next year only?" The answer to this would be compared to the theoretical value x solving $u((1+x)s)+e^{-\delta}u((1+x)s) = u(s)+e^{-\delta}u(1.1s)$.

As an example, suppose this year's budget is \$10 million and that a $50-50$ chance of no increase or an increase to \$11 million is just as attractive to us as having a budget of \$10.3 million for sure. From Table 4, we know that $u(s) = s^{1-R}/(1-R)$, the "felicity function" that indicates how we feel about different spending levels, and therefore relative risk aversion can be computed by solving

$$\frac{10.3^{1-R}}{1-R} = .5\frac{10^{1-R}}{1-R} + .5\frac{11^{1-R}}{1-R}$$

for R, which can be done by trial and error or by plotting the difference between the two sides as a function of R. In this case, we find $R \approx .498$. Suppose further that we are indifferent between flat spending this year and next at \$10.45 million or having spending of \$10 million this year and \$11 million next year. (We are willing to make do with a little less if we can do some of the spending now.) Then we can find δ as the solution to

$$\frac{10.45^{1-R}}{1-R} + e^{-\delta}\frac{10.45^{1-R}}{1-R} = \frac{10^{1-R}}{1-R} + e^{-\delta}\frac{11^{1-R}}{1-R} \quad ,$$

where we take $R = .498$ as already computed. Solving this (we can do so algebraically or numerically), we find $\delta \approx .177$. The optimal strategy is then given by Table 1 where r^* and K are given in the section "More Details of the Formal Model" in the Appendix. These formulas use the values of R and δ computed here, as well as values for the return parameters: real riskfree rate r, mean risky return μ, and standard deviation of risky return σ.

Having made the best case for this approach, I must say that I am a bit skeptical about the quality of the choices coming from such an approach. I don't think most people find these parameters very intuitive, and the paramaters may not be very useful if the assumptions about preferences are not accurate or if introspection about these examples is not very representative of our true preferences about actual results. For example, the derivation above may sound plausible, but the value of δ that comes out is probably unreasonably large, and the value of R is probably unreasonably small. The problem is that the sample questions used to elicit preferences are not very similar to realistic problems in endowment management. Furthermore, if we are not careful in selecting parameters, the solution may not exist at all. For example, if δ and R are too small, there may be more solution because any plan would be dominated by delaying spending to have even larger spending later. I prefer to explore preferences in the context of the choice problem at hand, for example using the simulation applet described above.

5 Further Implementation Issues

A formal model like the one we have described necessarily makes strong assumptions that are not true in practice. This section contains some suggestions on how to adapt the abstract model to the more complex practice.

Committed Spending

The theoretical model we have discussed so far takes the committed amount to be maintenance of current spending (probably in real terms) forever. One simple alternative that was considered in Dybvig [1995] is a commitment to decline at no more than a given rate. If the largest permissible decline is zero, then we have the model analyzed above. The decline in the commitment makes it possible to support a higher current spending level. Before, we could only spend the current endowment times the real interest rate; now, we can spend the current endowment times the real interest rate *plus* the rate of decline in the commitment. Of course, nothing comes for free, and the larger the possible rate of decline, the weaker the protection against lean times. As the permitted decline increases, the strategy looks more and more like the traditional strategy.

The revised strategy when the promise can decline at a rate d is given in Table 5. The portfolio strategy is different because the protected part of the endowment is smaller: the investment required to maintain current spending forever is less. This is pay for our spending not just from the riskless return but also some of the original investment. The consumption strategy is different because consumption declines at a rate d when we are not at the minimum proportion r^*. It is further different in a more subtle way since the minimum spending rate r^* may be larger than r (so long as it is less then $r + d$).

More generally, we can think of the commitment being any shape we might choose. For example, we might view ourselves as having a 5-year horizon and our commitment extends into year five. Next year, we will make a

Table 5: The Proposed Strategy with a Maximum Decline

Portfolio choice: The part of the endowment placed in risky assets is

$$\alpha_t = K(w_t - s_t/(r+d)),$$

and the amount placed in riskless assets is the remainder

$$w_t - \alpha_t = s_t/(r+d) + (1-K)(w_t - s_t/(r+d)).$$

Spending rule: Spending normally declines by a proportion d per unit time and is increased only as is necessary to maintain

$$s_t \geq r^* w,$$

with $r^* < r+d$, where r is the riskless rate of interest and d is the permitted rate of decline. That is, there is a critical proportion r^* of wealth below which spending can never fall. The plots in Figures 5 and 6 take $K = 1$ and $r^* = 1.5\%$/year.

commitment for what now looks like year six, and that commitment may be larger or smaller than the commitment for year five. We might also adjust the commitments in years 1-5. Of course, if we permit ourselves the freedom of reducing the commitments to years 1-5 we are making a bit of a mockery of the process, since that makes it unclear what if anything the commitment means.

Investing the Committed Part and Uncertain Interest Rates

Given a fixed interest rate, it is easy to figure how much we need to keep invested to honor our future commitment. If we commit to \$500,000 per year forever, and the interest rate is 5%, then we require \$10 million in the riskless asset to honor the commitment, since the annual interest on \$10 million is 5% of \$10 million or \$500,000. For more complicated cash flow streams (for example commitments of various amounts in the next 5 years), we can use the well-known concept of present value to compute how much we need to

invest to meet our commitments.[3]

In practice, we do not know what the interest rate will be in the future. Therefore, we cannot simply invest the money in short-term investments and be assured of the value we need later. If we are trying to maintain nominal (dollar) spending levels, there is a whole literature on *immunization* of portfolio value against interest rate shocks to ensure fixed payments. The simple approach is to match cash flows, although that is limited to maturities no greater than the longest available bond. More sophisticated approaches involve matching single or multiple duration measures, or better yet are based on a sensible multifactor model of the term structure of interest rates.

We are probably more interested in protecting the spending power of our commitments, which means immunizing in real terms. The good thing about protecting real spending power is that real interest rates seem to be much less volatile than nominal rates, reducing the reinvestment risk. Generally, there is a question of what investment can be considered riskless or nearly riskless in real terms, and one reasonable answer is that rolling over short Treasuries is pretty nearly riskless in real terms, especially over longer horizons. Now there is a good answer, namely that the new inflation-linked Treasury Bonds are riskless in real terms. For now, they are available for few maturities, which means that matching cash flows is impractical for now. However, the low volatility of real rates means that reinvestment risk is less important for these instruments, and we probably do not need many maturities for a good hedge. It will be interesting to see just how liquid this market is and how many maturities will be available.

It is interesting to note that there may be ways to convert real commitments to nominal. An example is the purchase of a building with borrowed money, which replaces the real need for space for offices and classrooms for a nominal commitment to make mortgage payments. This is related to another issue; buying the building is a more perfect hedge for that spending need than would be investment in indexed Treasuries, although its mortgage introduces

[3] If we have commitments $C_1, C_2, ...C_N$ at times $T_1, T_2, ...T_N$ and the interest rate is r, then the *present value*, which is the amount we have to invest at the rate r to meet the commitments, is given by $C_1/(1+r)_1^T + C_2/(1+r)_2^T + ... + C_N/(1+r)_N^T$.

inflation risk (since the payments are nominal) which can be hedged by guaranteeing nominal (rather than real) contributions of the endowment to the budget. In general we could think about hedging our demands for energy, salaries, and other expenses separately with different contracts, although our ability to hedge salaries accurately is probably limited.

Illiquid Assets

One potential advantage that endowment managers have over shorter-term investors is that they need less liquidity and can profit from longer-term investments. Illiquid assets do not pose a particular problem in the discretionary part of the portfolio, especially provided most of our portfolio is liquid and available for movement to the committed part as is called for. In the committed part, illiquid assets may be appropriate provided they give the required return pattern with little or no risk. This may, however, be uncommon among illiquid assets, because illiquidity is usually the result of information asymmetry about a risky asset.

Coordination with Other Sources of Income

Few educational institutions have endowments large enough to represent the majority of the budget. This observation leads to the question of how useful is it to have a guarantee on a small part of the budget. The answer to this probably depends on the reliability of other sources of income and on the use of the proceeds from endowment. If the proceeds to the endowment are earmarked for a specific purpose that represents a commitment, e.g. to funding a faculty chair, then it may make sense to treat this as committed spending and use the techniques in the paper. The other consideration is the stability of the other sources and uses of funds; having some guaranteed amount in the event of a shortfall elsewhere may be prudent. This is not a consideration if the endowment is truly small: if spending from endowment goes into the a general budget in which it represents 1% of spending that is moving around significantly from year to year, whatever is done with the

endowment is not going to have a huge impact (unless, of course, there are important expenses that cannot be funded out of your usual budget, say, because of a legislative restriction on how a state school can use tuition and funds from the state).

In terms of thinking about guarantees, new money from operations or from donations may make it reasonable to permit spending at a rate higher than the real interest rate and still expect never to reduce real spending. This approach is potentially a bit dangerous, since we do not really have any guarantee that donations will keep coming in at the same rate.

Transitional Issues

One interesting question is the optimal strategy for making a transition from a traditional program to a program of the sort described here. The most difficult part of the transition is that to maintain a real guarantee we may need a lower spending rate than we inherit. There are a number of reasonable ways to make the transition. For example, we can switch a portion of the money, or, if we are fortunate enough to have significant new donations adding to endowment, then we can manage the new part in this way.

6 Conclusion

I have proposed a new approach to endowment management in which the asset allocation and the spending rule are linked through the concept of committed expenditure. I hope it proves useful in helping educational institutions to prosper.

Appendix

More details about the backtesting

The backtesting is based on the monthly SBBI (stocks-bonds-bills-inflation) data provided by the Ibbottson Associates and distributed by CRSP (the Center for the Study of Security Prices at the University of Chicago). Each simulation is based on two asset classes, stocks and bonds. For the stocks, we use the S&P 500, which is an index of 500 stocks having large market capitalization. For the bonds, we use short US Treasuries with an average duration of 3-months. Real returns are used for all simulations, so all values are in terms of equal purchasing power. To compute real (inflation-adjusted) returns, we use the consumer price index (CPI). There is some concern about what is the right price index to use for what purpose, and there is also some concern about potential biases in the CPI. However, for us, performing the adjustment is certainly better than not adjusting for inflation, and should give us reasonable results.

Here is how real returns are computed. Let y be the ordinary return on an asset in a month, which is defined to be the change in dollar value over the month (inclusive of any coupons, dividends, splits, rights, etc.) divided by the dollar value at the start of the month. Further let i be the inflation rate over the month, which is the change in the CPI index value over the month divided by the value at the start of the month. Then the real (or constant-dollar or inflation-adjusted) return on the asset over the month is given by $Y = (1+y)/(1+i) - 1$, which is the same as the change in purchasing power divided by the purchasing power at the start of the month. The quantities y, i, and Y are often expressed as percentages, and may be annualized by multiplying by 12.

To operate a the new type of strategy with protected real spending, we need to be able to make investments that are riskless in real terms. That is possible now, thanks to the existence of the new indexed Treasury Bonds, but that was not possible during our sample period for backtesting. Making do with what data are available, I am using the SBBI short Treasury return as a

proxy real riskless asset, since over time its returns tend to move up and down with inflation. For computing the strategy, we also need to know how to evaluate future real cash flows. For this purpose, I assume a fixed real rate of 3.5%. This strikes me as a bit high (and is higher than the average real return on short Treasuries), but it is consistent with current market rates on the new indexed Treasuries, and I think the simulation results are reasonable. All the backtesting uses monthly data and the strategy is updated monthly.

For producing the figures, the value of the endowment is computed monthly, but spending is added up to a single figure for each year. This is consistent with the usual annual budget cycle, and avoids having distracting intra-year volatility in the graph that has nothing to do with reality. Variants of the traditional spending rule underly Figures 1 and 3. For the corresponding Figures 2 and 4, the portfolio rule dictates dividing the invested funds into equal proportions of stocks (S&P index portfolio) and bonds (short US Treasuries). The spending rule in Figure 1 is to spend 4.5%/12 of the portfolio value each month, that is, 4.5% annually. In Figure 3, the spending each month is 4.5%/12 of a 5-year average of the 60 most recent months' real portfolio values.

Figures 5 and 6 are based on the proposed new type of strategy. Under this strategy, spending is increased as necessary (and only as necessary) to keep it from falling below 1.5%/12 per month, and spending is never decreased. For the portfolio choice in the proposed strategy, we have a protected part that is invested 100% in short Treasuries and the remainder is invested 100% in stocks. Based on a 3.5% annual real rate (as discussed above), the protected part is taken to be monthly spending divided by 3.5%/12.

Figures 7 and 8 are based on the "diluted" version of the new type of strategy with a maximum rate of decline of 5% per year. Initial spending and minimum spending are both set at 4.5% per year. For the portfolio choice in the proposed strategy, we have a protected part that is invested 100% in short Treasuries and the remainder is invested 133.33% in stocks and short 33.33% in short Treasuries. The cushion is taken to be spending divided by 12%.

More details of the formal model

The constant K in the solution to Problem 1 as given in Table 2 is

$$K = \frac{\mu - r}{R^* \sigma^2}$$

where

$$R^* = \frac{\sqrt{(\delta + \kappa - r)^2 + 4r\kappa} - (\delta + \kappa - r)}{2r}$$

is a number between 0 and 1 and

$$\kappa = \frac{(\mu - r)^2}{2\sigma^2}$$

The constant r^* in the solution to Problem 1 as given in Table 1 is

$$r^* = r\frac{R - R^*}{R}.$$

References

Black, Fischer, and André Perold, 1992. "Theory of Constant Proportions Portfolio Insurance," *Journal of Economic Dynamics and Control* vol. 16, no. 3-4 (July-October): pp. 403–426.

Commonfund, 1996. *A Common Vision: Working in Partnership for the Benefit of All.*

Dybvig, Philip H., 1988. "Inefficient Dynamic Portfolio Strategies, or How to Throw Away a Million Dollars in the Stock Market," *Review of Financial Studies* vol. 1, no. 1 (Spring 1988): 67–88.

Dybvig, Philip H., 1995. "Duesenberry's Ratcheting of Consumption: Optimal Dynamic Consumption and Investment Given Intolerance for any Decline in Standard of Living" *Review of Economic Studies* **62**, 287–313.

Dybvig, Philip H., and Steven Heston, in preparation, "Preservation of Capital on Average is Not Good Enough."

Jorion, Phillippe, and William Goetzman, forthcoming, "Global Stock Markets in the Twentieth Century," *Journal of Finance*.

Keane, George F., 1996. "A Brief History: Dream Fulfilled," in Commonfund, 1996, *A Common Vision: Working in Partnership for the Benefit of All*.

Merton, Robert C., 1971. "Optimal Consumption and Portfolio Rules in a Continuous-time Model," *Journal of Economic Theory* vol. 3, pp. 373–413.

Murray, Roger F., 1996. "The Formative Years: A Founder Reflects," in Commonfund, 1996, *A Common Vision: Working in Partnership for the Benefit of All*.