Comments Welcome

The Theory of Endowment Management and Dynamic Portfolio Theory

by

Chester S. Spatt*

December 1999

*Mellon Bank Professor of Finance, Graduate School of Industrial Administration, Carnegie Mellon University. The author gratefully acknowledges extensive stimulating interactions with other members of Carnegie Mellon University's Investment Committee. However, the views expressed in this paper are solely the responsibility of the author and do not reflect those of the Carnegie Mellon Investment Committee or any of its current or past members.

1. Introduction

The main issues confronting university (and foundation) endowment management concern asset allocation and spending policies. The centrality of these issues is illustrated by the variation in relative growth of various endowments. Those universities (and foundations) that had large exposure to equity (and especially to domestic large capitalization or technology equity) in the 1990s had endowments that fared much more strongly than those with more modest exposure. The current spending decision, which reflects the tradeoff being made between current and future consumption, also greatly influences the future growth of an endowment.

The theory of optimal portfolio and consumption behavior of investors (e.g., Merton [1969, 1971]) offers many insights about asset allocation policies and consumption decisions for university endowment funds. In this paper we explore the implications of the theory for endowment management. [Note that Merton [1991] discusses optimal portfolio and consumption theory in the context of endowments, but emphasizes the interpretation of the model, rather than critiquing its limitations. For example, Merton [1991] emphasizes the potential impact of the riskiness and valuation of a university's gift stream and other cash flows upon the asset allocation decision for the existing endowment. Asset allocation

for an endowment should reflect the valuation and risk of the non-endowment wealth of a university in addition to its endowment.] Analogously, the nature of the actual endowment management problem and observed endowment management practices shed considerable light on the limitations of the traditional optimal portfolio and consumption models. Consequently, we highlight some of the essential differences between the decision problem confronting a university endowment and that facing (small) individual investors.

While observed spending rules typically are based upon smoothing of past fund values, our analysis suggests that such rules may not efficiently smooth the innovations to the value of the endowment fund looking forward. Instead, our analysis suggests that the current value of the endowment fund should drive the spending decision unless the endowment's preferences embody extremely large adjustment costs to changing spending. It is interesting that to the extent that observed spending rules are typically based upon three-year averages of past fund values, often the spending decision is not heavily linked to the current fund value as compared to the values from several years earlier. Even if a fund manager wishes to weight past fund values in his spending decision, the low relative weighting on the current values needs to be evaluated.

We extend this analysis to address the endowment manager's

asset allocation policy. Our analysis suggests that the allocation of the fund's assets should be determined from the current value of the fund rather than past values. This suggests the benefit to stable asset allocation policies rather than derivative-based securities in which the current asset allocation is linked to prior values of the endowment fund. The link to prior values of the endowment fund arises in both traditional portfolio insurance strategies in which the endowment manager reduces his exposure to risk after poor returns to limit his downside exposure and a "costless collars" program in which the manager is willing to truncate his upside returns for downside protection. The variation in effective exposure imposes considerable risk on the endowment fund and is inefficient relative to a stable asset allocation profile. An additional weakness of these derivative strategies is that because these do not reflect the structure of actual risks in the economy that need to be borne by some investors (i.e., the derivative strategies are simply "side bets" that are in zero aggregate net supply), the risk from variation in the effective asset allocation exposure will not provide the investor a sufficient risk premium due to the nature of the equilibrium pricing of risks that are not positively supplied in the aggregate economy.

Our paper also introduces the importance of strategic interaction among endowments in setting asset allocation policy.

We recognize that the operating costs experienced by many endowment funds and their ability to realize their underlying operating objectives will be negatively impacted by the financial success of their rivals. In particular, the risk faced by underendowed universities largely concern their relative performance in good states of the economy in that they are unable to retain their relative position vis-à-vis the larger endowment, which take on significantly larger dollar exposures to risk. Our analysis emphasizes the importance for asset allocation across endowment funds of the relative size of an endowment fund to the organization's underlying operating budget and the proportion of the operating budget funded by the endowment draw.

We organize our presentation as follows. In Section 2 we review the traditional expected-utility formulation and its implications for asset allocation and endowment spending choices. Then, we discuss the choice of spending rules by endowments and the linkage to the theoretical paradigm in Section 3. Next, we examine the pitfalls associated with various derivative strategies for asset allocation by endowment management in Section 4. The unique aspects of endowment management, especially the strategic interaction among endowments in asset allocation are explored in Section 5. We offer some brief concluding remarks in Section 6.

2 The Traditional Expected-Utility Problem: Consumption--Portfolio Theory Implications

Perhaps the most important issue for endowment management over time is the long-run risk exposure of its financial assets. The risk exposure issue addresses the broad risk profile (asset allocation) of the portfolio (the mix between equity and debt) and the specific allocation within the broad equity and debt universe to different investment styles. As a perpetual institution, the preferences of a university endowment can be summarized in terms of a long-run criterion function such as maximizing the expected utility from its operations over time. One can define the utility as being over its overall intertemporal stream of (real) expenditures or alternatively, the specific portion of its overall stream of expenditures that are generated by withdrawals of endowment income. While the former directly embodies how changes in the spending and investment policies can be used to hedge changes in the remaining real cash flows of the institution, the latter one captures and focuses upon the direct impact of differences in investment management policy.

The overall asset allocation and spending problem is most naturally formulated as a "consumption--portfolio problem" in which the dynamic stream of consumption withdrawals from the investment fund and the dynamic portfolio choices (among different

risky securities) are determined so as to maximize the institution's intertemporal criterion function. Versions of this problem, which have been tackled by scholars in economics, finance and decision theory, provide potentially useful qualitative guidance on a number of features of the optimal asset/risk allocation and the form of the optimal spending policy rule.

The problem confronting university endowments can be cast as maximizing the expected utility of the university's intertemporal consumption. For example, if c(t) denotes the endowment draw at time t, then the utility of the expenditure in period t can be computed from a utility function U(c(t)), where U(.) is an increasing and concave (i.e., risk-averse) representation of utility. [Alternatively, the utility could be defined over the entire university's expenditure, i.e., the expenditure from other sources as well as from endowment income.] The overall intertemporal utility function is then the sum of the discounted values of the utilities from the individual periods (such as by discounting at a constant rate). This constitutes the criterion for evaluating the university's preferences in the problem. Note that in this specification there is no direct utility impact of an increase in consumption expenditure in a period upon the expenditures in other periods (of course, there is an indirect effect through the budget constraint defined by the size of the endowment fund). Utility is separable across time periods and

states of economy in the traditional expected-utility formulation.

This preference relationship is maximized over the intertemporal portfolio allocation and expenditure decision from the endowment (of course, the more that is spent in an individual period the less that is available for the future). The solution to this problem has a number of standard theoretical features. For example, the heavier the discounting of the future the greater the current endowment expenditure.

A basic feature of the solution to the problem is that the optimal spending decision and asset allocation is a function of current wealth, but not the past wealth of the endowment or prior spending decisions. [This emerges from the first-order conditions to the investor's optimization problem and is a consequence of the separability in the expected-utility specification.] In fact, many observed asset allocations and spending rules do not satisfy this property. This issue is explored in more detail in the following section. Also for an important class of preferences (constant relative risk aversion, i.e., power or log utility) the asset allocation across assets and spending decision simplifies further. Under this parametric assumption both the optimal current spending and asset allocation is a constant proportion of the investor's current wealth. Indeed, many universities (and foundations) have adopted asset allocation and spending policies that are

proportional to the current or recent size of the endowment.

3. Optimal Spending Rules and Portfolio Allocations and the Dependence Upon Prior Wealth

The solution to the optimal intertemporal portfolio allocation and spending decisions should be forward looking. Under the assumption of zero transaction costs, past decisions and endowment performance should not affect the optimal current endowment spending decision and current portfolio allocation except through their impact upon the overall current fund value in an endowment (without the assumption of zero transaction costs, past portfolio decisions can have a slight influence upon the new optimal portfolio allocation). This is a direct consequence of formulating at each point in time the choice problem as an expected-utility optimization problem (looking forward) subject to the constraint on opportunities imposed by the current endowment and other wealth. In particular, this suggests that the level of endowment spending and the portfolio allocation at a point in time only depends upon past returns and past decisions through their impact upon the wealth of the endowment at that point in time. One implication of this is that past decisions should not influence the optimal portfolio mix going forward except through the effect on the current endowment level. Another implication is that while the optimal spending policy should depend upon the current

endowment it should not depend upon past levels of the endowment within this expected-utility framework. This argument suggests that a spending rule in which endowment spending (or asset allocation) depends in part upon the level of the past peak of the endowment is not optimal within an expected-utility maximization framework. It also suggests that the policy rule applied by many universities of determining endowment spending based upon an average of the level of the endowment in recent years is suboptimal in an expected-utility formulation. [Note that in the case of foundations Internal Revenue Service guidelines effectively require that the foundation contribute at least five percent of its current value to charities each year, subject to carry forward and carry back exceptions. This does not directly restrict the form of the spending rule, except in those instances in which the IRS standard is being violated.]

The conclusion that spending should be based upon the contemporaneous value of the endowment rather than prior values may seem surprising in light of the assumed aversion to risk and the benefits of smoothing by an endowment. While it is certainly true that the optimal policy of a risk-averse entity (I assume U(.) is a concave function) should embody smoothing of spending, the optimal rule should not be backward looking (i.e., it should not depend upon past outcomes or decisions) in an expected-utility framework, except through its dependence upon the contemporaneous

level of wealth. Instead, the sense in which the optimal rule does reflect "smoothing" is that innovations to the level of the endowment from gifts and returns are gradually spent over the future (incorporating "smoothing" in a forward-looking fashion). Notice that these implications are relatively general in that they do not require particular functional or parametric forms for the utility function U(.) within an expected-utility maximization framework.

In fact, academic researchers in recent years have begun to explore alternative formulations of the optimal consumption and portfolio problem (e.g., by relaxing the separability implicit in the traditional expected-utility formulation) in which past (as well as present) wealth influences the investor's decisions. For example, the typical form of a university's endowment spending rule is for it to be based upon the average value of the endowment over several years. In many of these cases the spending rule of the endowment is based upon the average value of the fund in the last three years or twelve quarters.

One possible interpretation of these spending rules is that these represent suboptimal policy decision rules. An alternative perspective is that casting the problem in a traditional expected-utility framework is misspecified and that under suitable alternative formulations past endowment levels would influence

current spending if preferences over spending reflected habit or relative spending effects rather than just the level of spending in each period. [The theoretical features and implications of such preferences are explored in Ryder and Heal (1973), Sundaresan (1989) and Constantinidies (1990).] Of course, even in such cases it might make sense to place the greatest weight on the current endowment level, a characteristic which few of the actual spending rules possess.

From a methodological perspective the endowment spending context is an interesting one for studying the dynamic consumption decisions because the spending rules of endowments are relatively straightforward to observe and indeed, in many instances reported by the institutions. This contrasts with the consumption policy of a traditional consumer, which would need to be inferred from demand or portfolio data. In light of the substantial correlation in wealth over time, it may be hard to estimate such policies for traditional consumer--investors, unlike endowments. The observed endowment spending policies can be used to help better understand the endowment's objective. The dependence of current consumption spending by endowments upon prior wealth suggests that habbit-like effects may be rather important for the endowments (and perhaps for the population more generally). Of course, one specific reason that habit effects may be particularly important for university endowments is the long-run nature of many of the commitments being

made (e.g., in the form of tenured faculty members and facilities) and the costliness of adjustment. These endowments seem to prefer a smoother consumption profile than the relatively volatile value of their underlying funds (note that the relative volatility of consumption compared to asset values motivated the preference specification in Sundaresan [1989]). Nevertheless, the widespread use of these spending rules by university endowments points to an important exception to the particular separability assumptions underlying traditional expected-utility models. Indeed, whatever the underlying cause, it seems efficient for universities to use smoother spending rules rather than altering their investment policy substantially so as to induce a smoother spending profile.

4. Asset Allocation and Optimal Portfolio Theory

As discussed above, traditional portfolio theory suggests that the optimal asset allocation policy to be followed by endowments should depend upon the current value of the fund but not prior values, past return realizations or past aggressiveness of the asset allocation (assuming zero transaction costs).

Indeed, it is even ambiguous whether the optimal risk mix should depend upon the fund's current value. The proportion of the portfolio held in various assets is independent of the level of wealth in many standard risk-averse formulations of the expected-

utility optimization problem. For example, if the utility function U(.) exhibits "constant relative risk aversion" (which occur when U(.) is a power or log function), then the portfolio mix among assets is independent of the level of the endowment, i.e., that the expected-utility framework results in the relative or proportional demand for the risky assets not depending upon the endowment fund's value. This is a standard and important class of preferences that is widely employed in academic research. Under this assumption, an endowment's proportionate asset allocation would not be influenced by its fund value.

Endowments have used a variety of derivative strategies to alter the profile of their risk exposure. One common form is "portfolio insurance" in which an institution reduces its exposure to equity as the market value declines (and increases the exposure to equity as the market value increases) in order to create a downside floor under the portfolio's value to provide portfolio insurance (this is illustrated by the presentation in Grossman [1988]). Our discussion above concerning the constant relative risk-averse preferences suggests that "portfolio insurance" is not part of the optimal solution for many risk-averse investors. More generally, it is certainly unclear that a risk-averse investor should have his demand for risky assets increase more than proportionately as wealth rises.

Another popular type of derivative strategy is the "costless collars" program in which institutions obtain downside protection by buying out-of-the-money index puts that are financed by the sale of out-of-the-money index calls. There are several features of the "costless collars" program that merit attention:

a) Obviously, the "costless collars" program substantially reduces an endowment's exposure to equity (substantially below the stated level) and more specifically to the reference index. This reduction in equity exposure is beyond any conscious (explicit) underweighting.

b) More interestingly, the actual effective exposure to equity under the costless collars program is sensitive to the level of the reference index (stock price) in that the sensitivities of the values of the various puts and calls used in the program change with the reference index. More specifically, as equity values rise the sensitivity of the call option to the value of the underlying index increases (towards one) and the sensitivity of the put option declines (towards zero) because as the stock price increases the probability of the call being exercised goes toward one, while the probability of the put being exercised becomes almost zero. Because the endowment is buying puts and selling calls, its overall effective exposure is declining with the level of the index. This involves an unconscious "timing bet" relative

to a constant portfolio mix. If the market is unpredictable and the distribution of return opportunities in the market is stable, then the costless collar strategy would be dominated for risk-averse investors by a more stable asset allocation (this is a consequence of expected return being linear in the asset holdings, while the risk-bearing costs of increased risk exposure grow nonlinearly). Risk aversion per se does not imply that the costless collars approach is an efficient way to reduce risk. Similarly, it is unclear that the effective asset exposure should depend upon recent returns. The same critique would apply to the traditional portfolio insurance strategy described above, though there the implicit equity would be increasing (rather than decreasing) in the value of the underlying reference portfolio.

c) Because the effective exposure of the endowment is sensitive to the underlying level of the index it is important for endowment managers to monitor the actual sensitivity of the portfolio value to the value of the underlying index.

d) From a risk-bearing perspective it is unclear that there is compensation for bearing the variation imposed by a "derivatives" style strategy such as a "costless collars" or "portfolio insurance" program. In particular, these strategies do not reflect the (stable) bearing of risks that are positively supplied in the economy, but strategies that are in a zero aggregate (net) supply

and therefore, do not earn a "risk premium" in equilibrium.

e) As an aside, note that the implicit prices of payoffs in states in which the reference index is substantially below the current market level are artificially high relative to that for cash flows in more favorable states (this has been extensively documented by academics in studies of volatility "smiles" in recent years). This suggests that the put options being sold are very expensive (perhaps due to the artificial demand for such payoff patterns) to reduce equity exposure.

One of the interesting implications that emerges from an intertemporal portfolio framework is that it implicitly identifies the types of risk that matter for university endowments. As long horizon players, the crucial risks are those about long-term returns. Monthly variations in returns do not drive the evaluation of the performance of an endowment as measured by the long-run criterion function. Instead, one is interested in the sustained performance of the endowment (and individual managers) over long time intervals. Indeed, it is only annual or quarterly returns that even influence the permitted endowment spending draws. Even apart from the theoretical argument, as one looks back at the performance of an endowment over the last 10 (or 25) years, it is clear that we are inherently interested in the ability of the endowment to sustain and improve an institution's programs (both

absolutely and relative to its competitors) rather than monthly (or daily) variation in the fund's balance sheet.

5. Unique Aspects of Endowment Management and Implications for Asset Allocation

There are several important differences between the problem confronting an individual investor and an endowment. While some of these are differences in degree, some of the distinctions are relatively more fundamental.

One difference between an individual investor and an endowment is the perpetual nature of a university (or foundation). However, the "infinite horizon" model in which the investor's lifetime does not have a prespecified bound is an ideal way to capture this feature of the problem. Indeed, the policy (e.g., asset allocation and spending) decisions of an "infinite horizon" investor are not that different from those of a finite horizon investor whose horizon is relatively long. Of course, one of the inherent comparative advantages of the endowment investor is its long horizon and relatively modest potential needs for immediate liquidity. This provides the endowment investor greater flexibility (than individual investors) to hold relatively illiquid assets (though the endowment investor does not have an obvious comparative advantage in this dimension over investment by

a large pension plan that covers many individuals) for a reasonable portion of its portfolio. [Of course, among the limitations of this perspective are that (a) at periodic intervals the investment committee reports its performance and (b) the membership of the committee is for limited terms. In fact, these features can transform an endowment from following the objective of a true perpetual institution solving an infinite horizon problem to one whose horizon more closely resembles a typical individual investor.]

Another important feature of endowment management is its tax treatment. Endowments are essentially tax exempt. [Note that a foundation is subject to a 2% tax on its investment income, but the modest nature of this tax would not lead to a significant change in investment orientation.] This is a major comparative advantage, which enhances somewhat the investor's investment opportunity set and skews the investor towards ownership of more heavily-taxed assets. [However, in practice these more heavily-taxed assets are often fixed-income securities or high-yield equities that are not attractive to endowment funds with a long horizon.] In fact, if we ignore the horizon effect (which is itself rather important) the theory of tax clienteles suggests that endowment managers should own relatively more heavily-taxed securities (such as corporate bonds) as compared to ownership of those assets by heavily-taxed investors. Similarly, endowments

should strive to avoid UBIT (Unrelated Business Income Tax) on their investments by structuring the claims that they hold appropriately. Indeed, some transactions proposed for endowments have considerable implicit UBIT liability.

Competition Among Universities and the Ability to Take Risk

An important dimension of the problem confronting a university endowment is how does the competition that it faces influence the determination of the optimal spending decision and portfolio asset allocation. Perhaps the feature of the endowment management problem that is least well captured by the standard academic perspective concerns the impact upon the investor's utility of returns earned by other investors. For most investors, the principal impact of returns being earned by other investors is through the benchmarking and evaluation of one's own asset managers. In the case of university endowments, the effect of the performance of one's rivals is much more fundamental. Specifically, the competitive posture of a university and its ability to attract talent (outstanding faculty and students) is greatly influenced by the size of its resource base (both through the direct offer received by the individual and the indirect signal about the future attractiveness of the university to others). For example, the "cost" of high quality faculty members that a university is attempting to attract or retain would

increase with the investment performance of its rivals. [This point was explicitly emphasized in a 1998 New York Times account of Columbia University's effort to recruit economist Robert Barro from Harvard and the more general lifting of prior caps on the salaries of outstanding senior economists as a result of the successful performance of the endowments at many leading universities.] Consequently, endowment investors are concerned about their relative performance as well as their absolute performance. This discussion points to a central strategic aspect to the investment decision process. [While one could argue that such strategic effects are important for individual investors, e.g., if they are competing for a limited resource, such as beachfront retirement property in San Diego, that case seems more strained.] The standard optimal consumption--portfolio theory problem focuses upon an investor in isolation and does not directly address the types of strategic issues (in which the decisions of one's rivals interact with one's own utility) that are important to decision making for a university endowment.

University endowments are long horizon players in the financial markets and accordingly possess considerable exposure to equity and other relatively risky assets. The problem confronting a university with a relatively small endowment is that the relative size of its endowment is destined to shrink in "good" market states if it bears relatively less risk. But these "good"

states are particularly problematic for an underendowed university, even if it possesses the same relative risk exposure, because it is in these scenarios that the university with the relatively small endowment possesses the greatest competitive disadvantage (as it has more limited financial assets) and pressures on faculty compensation due to overall faculty labor market conditions and raids of faculty by competitors.

An interesting example of strategic aspects concerns the lack of equity exposure by New York University due to concerns about market "valuation" (based upon accounts in the Wall Street Journal in the fall of 1997). This was particularly problematic due to the great endowment success of some of NYU's competitors during that period. [However, NYU enjoyed considerable success in other dimensions during the 1980s and 1990s. This may have reflected (in part) implicit financial market exposure due to its New York City location.]

Another way to look at the strategic problem is that an important variable is the proportion of a university's operating budget that comes from endowment income. Universities whose endowments are large relative to their operating budget and are therefore able to have the endowment sustain a more substantial portion of its budget are relatively more exposed to market declines. A university receiving 23% of its budget from its

endowment receives a 4.6% annual budgetary decline from a 20% market hit, while another university receiving only 7% of its budget from its endowment feels a 1.4% decline in its annual budget from a 20% market hit (assuming identical investment policies and identical consumption spending rates of 5%). In this sense the wealthy university is actually more sensitive to changes in its endowment and less able to take proportional risks with its endowment (even if it possesses somewhat greater ability to defer maintenance, etc.). Of course, this does not suggest that a university with a higher endowment is worse off, but rather that the larger endowment actually leaves a university more vulnerable to market declines because the university with the larger endowment has spending commitments reflecting its larger financial base. The implication of these observations is that a university with a relatively smaller endowment should consider a relatively more aggressive (rather than less aggressive) investment posture. This is not captured in the standard consumption--portfolio framework, which treats the individual investor in isolation. Erosion in the relative size of an endowment can have strong negative consequences for the ability to attract high quality students and faculty and external gift and grant support. For example, a university with a relatively smaller endowment may need to contract or be unable to compete if its endowment grows relatively less quickly than its competitors. The financial base created by the endowment is important in attracting the financial

resources and intellectual assets that are central to the university.

6. Concluding Comments

The discussion in this paper highlights several issues in which the intertemporal framework described suggests surprising implications for spending and asset allocation policies. Among the interesting issues examined are the form of the optimal spending rule and the dependence of current vs. prior fund values, asset allocation and traditional derivative strategies such as portfolio insurance, and the importance of strategic consideration in asset allocation (e.g., should relatively smaller endowments bear proportionately smaller or larger risks than its competitors). The contribution of this paper is to identify limitations in traditional intertemporal consumption and portfolio theory vis-a-vis endowment management and provide fresh insights on a number of features of conventional decision rules.

References

Constantinidies, G., 1990, "Habit Formation: A Resolution of the Equity Premium Puzzle," Journal of Political Economy, 98, 519-43.

Grossman, S., 1988, "An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies," Journal of Business, 61, 275-298.

Merton, R., 1969, "Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case," Review of Economics and Statistics, 51, 247-257.

Merton, R., 1971, "Optimum Consumption and Portfolio Rules in a Continuous Time Model," Journal of Economic Theory, 3, 373-413.

Merton, R., 1991, "Optimal Investment Strategies for University Endowment Funds, Harvard Business School Working Paper #90-029.

Ryder, H. and Heal, G., 1973, "Optimum Growth with Intertemporally Dependent Preferences," Review of Economic Studies, 40, 1-33.

Sundaresan, S., 1989, "Intertemporally Dependent Preferences and the Volatility of Consumption and Wealth," Review of Financial Studies, 2, 75-89.