William Massy • Jackson Hole Higher Education Group

CAPITAL STRUCTURE and RISK MANAGEMENT

For decades investment managers at colleges and universities have worked to optimize endowment asset portfolios. The techniques for doing so have evolved from simple mean-variance analysis and efficient frontiers to increasingly sophisticated simulations of portfolio returns and their effects on spending. A variety of spending rules have been developed to filter the effects of investment return volatility while maintaining the endowment's purchasing power. William Massy, Professor Emeritus of Education and Business Administration at Stanford University and currently president of the Jackson Hole Higher Education Group, takes these efforts a step further by analyzing the portfolio of income and expense items outside the endowment, i.e., the operating portfolio. He presents a financial planning model to help managers think about the operating portfolio, both in and of itself and in relation to investment strategy, and to consider how major changes in operating-item volatility might be mitigated by countervailing changes in other operating items and/or asset allocations.

NOTEBOOK

- Given the intense focus on the financial management of endowments, it is perhaps surprising that relatively little attention has been devoted to managing risk in the operating portfolio and its relation to investment risk.
- Getting used to looking at risk profiles and developing methods for mitigating risk should be an important goal for university planners.
- University endowments potentially can be used to hedge the effects of operating risk.
- It is only by quantifying desired shifts in asset allocation targets and their effects on the university's optimal capital structure that the possibility of using the endowment to hedge operational risk can be put on the table.

Analyzing Operational and Investment Risk

Given the intense focus on the financial management of endowments, it is perhaps surprising that relatively little attention has been devoted to managing risk in the operating portfolio and its relation to investment risk. (For convenience I include debt service on the expense side of the operating portfolio even though it is not an operating item per se.) Items in the operating portfolio can be as volatile as the returns from some asset classes, and the fluctuations for the various items may well be correlated with each other and with the asset-class returns. How, then, might major changes in operating-item volatility be mitigated by countervailing changes in other operating items and/or asset allocations?

Suppose, for example, that a university embarks on a large and risky new research venture. The program will bring significant increases in direct and indirect research revenues and the costs associated with them. Construction of a large new laboratory also will be required, which will be funded by debt. Let us assume that the university has developed a plan that brings the expected values of the revenue and cost increases into balance. In other words, the program will work financially if everything comes out as planned.

But things almost surely will not come out as planned. First, the direct research-revenue stream will be subject to the vagaries of agency and corporate funding, gifts, and, perhaps, the ability of faculty to field competitive proposals. Most of the direct costs will vary with the direct revenues, but some will do so only after a lag or not at all. In the short term, the university may be required to provide substantial bridge funding to fill gaps between projects and perhaps support some programs for long periods. The indirect-cost side may be even stickier. For example, capacity put in place to serve the extra faculty and staff may not be shed easily if project revenue declines. This is particularly true for the operations and maintenance costs of the new laboratory. Finally, the indirect cost rate itself may not always yield full recovery, especially if the university operates under NIH cognizance.

It is not easy to quantify these uncertainties using conventional spreadsheet models. Yet their effect will be significant if the fluctuations are large compared to those associated with the university's normal operations. For example, the new research revenue may be considerably more volatile than existing research revenue streams and, for selective institutions, tuition revenue. Utility rate increases for the new energy-intensive research facility may introduce unexpected volatility. The same may be true for supplies and, perhaps, the salary rates for specialized faculty and staff needed to support the new program.

The downside effects of such volatility will have to be addressed by some combination of deficits (i.e., by tapping the university's operating reserve) and budget reductions or income enhancements in other areas. None of these will be painless, yet our knowledge and even our intuition about them are primitive. How, then, can the effects be analyzed and incorporated into the university's financial planning regimen? And even where the analytics aren't formally incorporated in planning models, how can we improve our subjective understanding of the effects of operating risk?

Visualizing Operational Uncertainty

The easiest method for analyzing uncertainty in the university's operating portfolio is similar to the one used to analyze the investment portfolio: that is to say, Monte Carlo simulation. One models the linkages among the revenue and cost items, makes assumptions about their standard deviations and correlations, and then calculates and displays a measure of volatility for, say, the university's operating margin. But while the approach is straightforward and well known to investment analysts, it is typically not part of a financial planner's toolkit. Hence, the analysis of uncertainty in the operating portfolio is something of a rarity.

My newly developed "comprehensive financial planning model" (CFPM) is designed to help planners analyze uncertainty in the operating portfolio, among other things. The model is being implemented at the National University of Singapore. Through the use of graphic displays, it helps planners visualize the consequences of forecasting and policy assumptions, including assumptions about volatility, revenues, expenses, and the changes in funds balances by funds group and organizational unit. The displays also illustrate what is almost a universal feature of probabilistic models: the tendency for the variance of an uncertain quantity to increase over time. Getting used to looking at risk profiles and developing methods for mitigating risk should be an important goal for university planners.

A Capital Structure Model for Nonprofit Enterprises

Capital structure models for profit-making ventures depend on the volatility as well as the expectation of profit. Volatility and the anticipated growth rate of profit determine a stock's price-earnings ratio, which, along with the current profit expectation, determines the stock price. Things are similar in the not-for-profit world, except that total spending (or spending per student) replaces profit as the criterion. The substitution derives from the fact that a school's accomplishments depend on its activity levels, which in turn depend on expenditures. The deleterious impact of volatility on price-earnings ratios is mirrored for nonprofits by the difficulty and pain of budget reductions, and of course the problem of bankruptcy is the same in the two cases.

The model used in this paper focuses on operating margin, the accounting analog to business profit, rather than on total expenditures. This strategy avoids the need to model an institution's decision rule for apportioning revenue and cost variations between budget adjustment and reserve usage. The difference between operating margin and total expenditures, or any linear decision rule for apportioning revenue and cost variations for that matter, is not important for present purposes.

It can be difficult to convince trustee investment committees to change asset allocation targets, but that shouldn't deter one from analyzing the potential advantages of such changes.

To build such a risk model for nonprofits, we start with the volatility of the enterprise's revenue and cost streams and the returns on its various assets—which drive the volatility of operating margin. Balance sheet considerations enter the model in two ways: first, because they determine the dollar values of investment return and debt service, and second, because they limit the downside excursions of margin that must be absorbed by the institution's reserves. Substantial reserves permit substantial deficits, at least for a while, and conversely.

Operating margin (*OM*) is represented by the following:

$$OM = \Sigma_I \, ORC_i - DS \, D_0 + p(1 + \Sigma_j \, x_j TR_j)E_0 + ONR, \qquad (1)$$

where

ORC_i represents the *i*'th operating revenue or cost item

p is the spending rate on the endowment

DS is debt service as a percentage of principal

D_0 is beginning debt principal

x_j is the fraction of endowment allocated to the *j*'th asset class

TR_j is the total return for the *j*'th asset class

E_0 is the beginning endowment value

ONR, other net revenue, is the net of all revenues and costs not included elsewhere in the model. *ONR* is not subject to uncertainty.

The basic equation could apply to any not-for-profit enterprise, but the introduction of endowment and debt service particularizes it to colleges and universities. The model reflects only one year of operations and ignores the smoothing of endowment returns, although it is possible to take smoothing into account.

The first two terms represent the operating portfolio and the last term represents the investment portfolio. The uncertain variables are ORC_i, TR_j, and possibly *DS*. The mean of operating margin is a linear function of the means of the variables in the two portfolios and the variance is a linear function of their combined variance-covariance matrix. Both the mean and variance depend linearly on the endowment asset allocations. Hence one can see the potentially offsetting or reinforcing effects of risky outcomes in the two portfolios.

The effect of risk on operating margin depends on the asset allocations, debt service, and the mix of income and expenditure items. The full-fledged capital structure model includes the size of the operating reserve, other reserves if applicable, and the details of the institution's debt covenants. For now, we simply look at the variance of operating margin. A large variance increases the chance that a reserve limit or debt covenant will be violated, and conversely.

Among other things, we consider the use of endowment as a hedge for operating risk. This would require asset allocation targets to change as a result of major shifts in the operating portfolio's risk profile, including those caused by changes in debt service. It also may be possible to design investment vehicles specifically for the purpose of hedging, though such ideas fall outside the scope of this brief.

It can be difficult to convince trustee investment committees to change asset allocation targets, but that shouldn't deter one from analyzing the potential advantages of such changes. After all, it wasn't so long ago that investment committees were reluctant to consider assets such as foreign securities, venture capital, private equities, and other classes that have become standard vehicles for many endowments. It is only by quantifying the desired shifts in targets and their effects on the university's optimal capital structure that the possibility of hedging operational risk can be put on the table.

The classic method for optimizing asset allocation begins with calculating the efficient frontier between risk and return using mean-variance analysis. That is, one minimizes portfolio variance subject to a succession of values for expected return. Then one evaluates the institution's desired trade-off between risk and return and finds the point on the efficient frontier where the frontier is tangent to the slope of the trade-off function. The optimum asset allocation is the one associated with this point. The method is easily extended to use operating margin instead of investment return. One can test the interactions between the operating and investment portfolios by changing the correlations among the elements of the two portfolios, calculating the new efficient frontier, and finding the point of tangency.

It is possible to simplify the method by using an explicit risk preference function—also called a "utility function." Then one can simply maximize utility instead of calculating the whole efficient frontier and finding the tangency point. David Swensen's book *Pioneering Portfolio Management* (2000) dismisses utility functions as impractical, but I'm convinced they can prove useful in the kind of analysis being conducted here. We shall use the following simple utility function:

$$Utility = a + b\,\{E[OM] - (1/k)\,\mathrm{Var}[OM]\}, \qquad (2)$$

where "E" means expectation, "Var" means variance, k is the model's risk aversion parameter, and a and b are normalizing constants set to make *Utility* =100 in our base case. Larger values of k mean a greater tolerance for risk. Because Equation (2) is quadratic, maximizing utility with respect to the asset allocations always produces a point on the efficient frontier. All that's necessary to determine k is for one's current asset allocation to be near-optimal given current assumptions; then simply maximize utility for different k-values until the results equal the current asset allocations. This simple method of determining k is what makes the utility model practical.

Mean-variance analysis tends to be unstable, so most investment managers now favor simulation approaches for determining asset allocations. However, preliminary tests of capital structure models can be conducted quickly and conveniently using the mean-variance utility model and refined later using simulation.

Application of the Model: "Virtual Endowment" Example

This example is based on my experience as Stanford's CFO, where I asked myself whether the university's expectations for large gifts from venture capitalists should affect the fraction of the endowment invested in venture capital. I had no tools for investigating the question and so let the matter drop, but now it is possible to do the analysis. Let us characterize those large future gifts as coming from a "virtual endowment." We assume an annual flow of $20 million, a standard deviation of $2 million, and estimated correlations with asset class returns of .50 for private equity, .30 for foreign equity, .50 for domestic equity, .35 for real estate, and .20 for bonds. This kind of volatility arises because many potential donors are heavily invested in venture capital (or, as in Stanford's case, are venture capitalists themselves) and thus are more likely to give when venture capital is doing well.

One might reasonably ask how one can know the correlations between gift flows and investment returns. The answer, of course, is that one can't really know them. However, an established tenet of management science holds that it is better to use one's best judgment about an unknown factor than to leave it out of the model and thus implicitly assume it to be zero—where zero is known to be wrong. A good approach is to obtain the best judgments of knowledgeable people, either by direct questioning or through a more elaborate technique such as the Delphi method.

Figure 1 compares the operating margin efficient frontiers for virtual endowment with and without correlations, as compared to the base case of a gift stream without risk—that is, a fixed gift stream with no variation. Adding virtual endowment *without* correlations shifts the efficient frontier slightly to the right, as would be expected given the increased risk in the system. However, the big shift to the right comes *with* correlations: Because the gift correlations are positive, they amplify the effects of investment return volatility when they are taken into account. This is just what one would expect if donors are responding to the same financial market factors as those that drive the endowment.

The resulting asset allocations and statistics for operating margin are shown in Table 1.

The first column of data in Table 1 shows our base case, in which there is no variation in the gift stream and, therefore, no correlations. Asset allocations were determined by maximizing the utility function. The result, which is equivalent to the standard asset allocation procedure described earlier, falls on Figure 1's "base case" efficient frontier.

Figure 1. Efficient Frontiers for Gifts to Operations (with and without correlations)

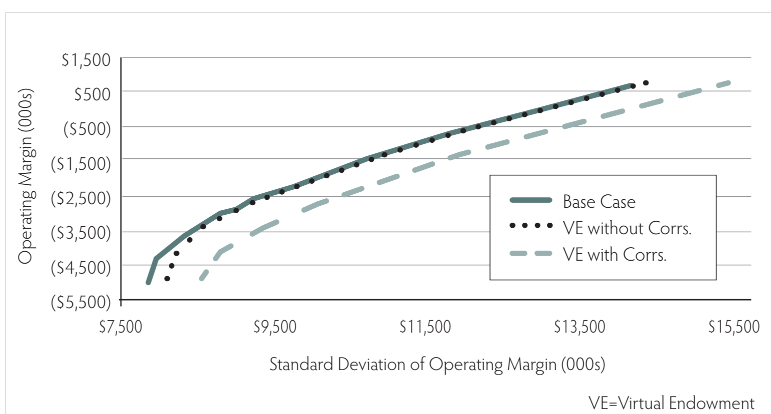

Table 1. Effects of Virtual Endowment on Portfolio Statistics and Asset Allocations

No Asset Class Adjustment	Base Case		No Correlations		With Correlations
expected utility	100.0		98.2		83.3
expected operating margin	$0		$0		$0
std. dev. of operating margin	$13,545	STEP 1A	$13,692	STEP 2A	$14,866
expected total return	7.28%		7.28%		7.28%
std. dev. of expected total return	13.5%		13.5%		13.5%
With Asset Class Adjustment					
expected utility			98.2		83.9
expected operating margin			$0		($479)
std. dev. of operating margin			$13,692		$13,981
expected total return			7.28%		6.80%
std. dev. of expected total return			13.5%		12.7%
		STEP 1B		STEP 2B	
Asset Class Allocations					
private equity	28.1%		28.1%		25.7%
foreign equity	21.1%		21.1%		20.8%
domestic equity	27.1%		27.1%		22.7%
real estate	13.3%		13.3%		11.0%
bonds	10.4%		10.4%		19.9%

Note: Dollar amounts are in thousands.

The following steps use the model to analyze the effects of operating risk:

Step 1, the middle column of the table, recognizes the variation in gift flows but ignores their correlation with asset class returns.

Step 1A shows the effect of these variations when asset allocations are held equal to those of the base case. Expected utility has declined from 100 to 98.2 because the curved utility function penalizes downward excursions more than it rewards gains. Expected operating margin remains at its base case level but its standard deviation increases from $13,545 to $13,692. Expected total return and its standard deviation are unchanged because the asset allocations equal those of the base case.

Step 1B shows the effect of maximizing utility with respect to the asset allocations. The answer is that nothing happens. Uncorrelated variation in gifts—and by extension all operating revenue or cost items that, however correlated they may be among themselves, are uncorrelated with asset class returns—does not affect asset allocation.

Step 2, the table's right-most column, recognizes the correlations as well as the variances. This, of course, is the more realistic case.

Step 2A shows expected utility at 83.3, a drop of nearly 17% from the base case. The standard deviation of operating margin increases by nearly 10% to $14,866. These substantial changes are caused by the amplifying effect of the correlations—that is, because the variations in gift flows are now more in synch with those of total return. Once again, the statistics for total return (7.28% expected total return and 13.5% standard deviation) are unchanged because the asset allocations haven't changed.

Step 2B shows how changes in asset allocation can mitigate the effects of correlated variations in operating revenue. Maximum utility is obtained with the allocations at the bottom-right of the table, which reflect a substantial increase in fixed income securities and a corresponding reduction in equities. This produces a lower total return than the base case and, consequently, a negative expected operating margin. However, the loss of $479 (thousand) is more than compensated for in utility terms by the reduction of the standard deviation of operating margin by $885 (thousand), from $14,866 to $13,981.

The example demonstrates clearly that an institution's optimal asset allocation can be quite sensitive to variations in operating revenue and expense—provided these variations are correlated with asset class returns. As for asset class returns, the revenue and expense variations will depend significantly on general economic factors such as inflation. Hence, such correlations are more likely to be the rule than the exception.

Refining and Extending the Model

To this point we have focused on the mean and variance of operating margin and the ways they may affect utility and asset allocation. For these to be sufficient statistics, however, requires two strong and not necessarily appropriate assumptions:

1. The variables must follow a multivariate normal distribution. Contrary to this assumption, investment returns may be distributed log-normally or exhibit what has come to be called "tail risk"—where the associations among variables are highly nonlinear and the probability of strongly negative returns exceeds that expected with normality.

2. There can be no limits or asymmetries in the university's ability to cope with risk. This assumption will be violated if the size of the university's operating reserve, or any debt covenants based on operational results or asset balances, limit the negative excursion of operating margin.

Such matters lie at the core of nonprofit capital structure optimization. They are discussed in detail in the full-length version of this paper, which will be made available at www.educause.edu/forum.

Further, the model for operational and investment risk can be extended in a number of important ways. One way is to consider multiple time periods. Another is to incorporate a smoothing rule for spending from endowment. Still others involve asymmetric probability distributions and those with nonlinear associations among the variables. The implementation of such extensions is beyond the scope of this summary, but my full-length paper discusses them in some depth.

William Massy is Professor Emeritus of Education and Business Administration at Stanford and president of The Jackson Hole Higher Education Group, Inc. He co-authored *Planning Models for Colleges and Universities* (1981) with David Hopkins. His most recent book is *Honoring the Trust: Quality and Cost Containment in Higher Education* (2003). Massy can be reached at bill@jhheg.com.

John Core • The Wharton School, University of Pennsylvania

DISCUSSION

Massy's points are at once both obvious and revolutionary. His basic idea is to manage not only endowment portfolio risk but to manage operating portfolio risk as well, and to use the endowment portfolio to manage that operating portfolio risk. It seems obvious (once Massy makes the point) that university managers should think about the risk of their total portfolio and not simply the risks of the operating portfolio or the endowment portfolio separately. Yet, this concept is revolutionary in the sense that for a number of reasons—many of them good—managers have been taught to consider the endowment as a separate entity that generally should be left untouched.

Using the Endowment as a Risk Management Tool

A revolution in risk management has occurred during the past 20 years or so. Stulz (1996) provides an excellent summary of the changes in risk management theory during that time, but he focuses his discussion on public for-profits. The basic idea of risk management is to make a firm's cash flows less volatile. Often this is done by "hedging." For example, a gold company may buy futures on gold prices. When gold prices fall, the futures become more valuable, so the decline in the gold company's cash flow due to falling gold prices is offset by the increase in value of the futures position.

Figure 1 (Stulz 1996) shows the distribution of firm value before hedging (dotted blue curve) and after hedging (solid blue curve). On one hand, it is puzzling that public for-profits hedge—since their equity is held broadly by many shareholders who can diversify risk; indeed, it is argued that for-profit firms should be run as if they were risk neutral and therefore do no hedging. But as Stulz discusses, deadweight costs of bankruptcy change this argument. If a firm violates debt covenants (e.g., its market value falls below the book value of debt), in an extreme case the debtors can seize the firm, limit operations, and destroy equity value. Since these bankruptcy