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 if desired.

Perhaps surprisingly, however, little progress has been made in analyzing the other main portfolio that falls within the purview of financial planners: the portfolio of income and expense items outside the endowment (the "operating portfolio"). This paper focuses on risk in the operating portfolio and its relation to investment risk. For convenience I will include debt service on the expense side of the operating portfolio even though it is not an operating item per se. This positioning should not suggest that debt service is "just another cost," however: The consequences of debt service being contractually fixed will be addressed later under the rubric of "downside risk."

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, should we think about the operating portfolio, both in and of itself and in relation to investment strategy? How 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 of all, 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

William F. Massy
Jackson Hole Higher Education Group

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 taken up 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 if the research-driven fluctuations are large, 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?

This paper presents new tools for the simultaneous analysis of operational and investment risk together with worked examples of how the tools might be used. But because such usage requires something of a paradigm shift, it is too much to expect the examples to be fully realistic or even fully representative of the problems where the tools may eventually be applied. This is entirely consistent with the goals of the Forum's Master Class, which has developed a robust tradition of proposing paradigm shifts. We hope that the discussions will unearth additional examples and illuminate areas for further study.

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 and indeed is described at length in some of my earlier work (Grinold, Hopkins, and Massy, 1978; Hopkins and Massy, 1981; Massy, Grinold, Hopkins, and Gerson, 1981), it is typically not part of a financial planner's normal 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

Figure 1. Uncertainty Profiles for Operating Margin

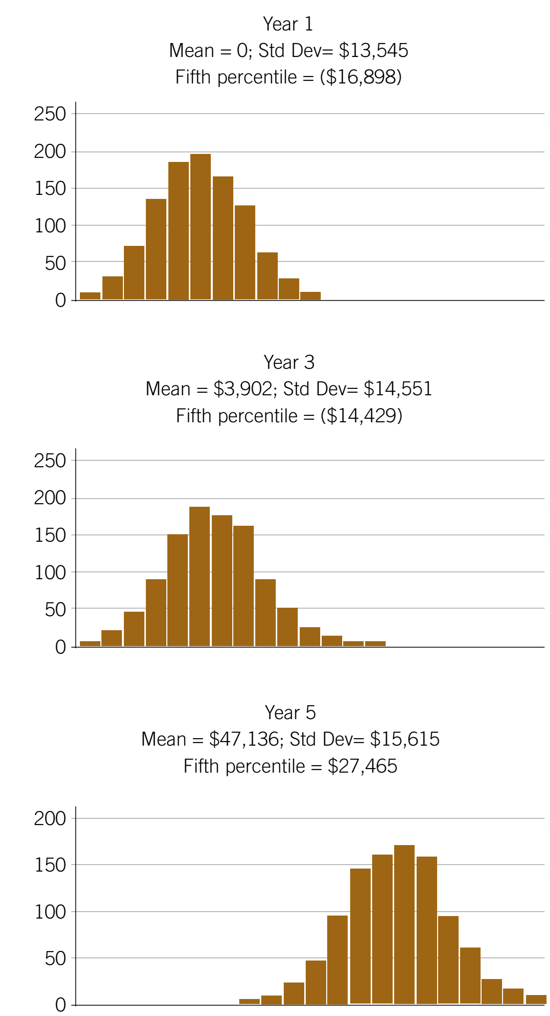

things. The model is being implemented at the National University of Singapore. It helps planners visualize the consequences of forecasting and policy assumptions, including assumptions about volatility, revenues, expenses, and the changes in funds balances by funds group and organizational unit. CFPM provides an Excel-based platform for illustrating the effects of uncertainty in the example described above.

Figure 1 presents the "bottom line" risk analysis from CFPM. The model has been configured to illustrate a hypothetical though not atypical university. To analyze the example, one simply enters the input data and double-clicks on the line for which the uncertainty profile is desired—in this case the operating margin or surplus (deficit). The uncertainty profile is represented by mean, standard deviation, frequency distribution, and fifth percentile of the frequency distribution for years 1, 3, and 5 of the planning period. (The years can be modified with simple program changes.) The random simulation is based on 1,000 trials, a number that can be changed by accessing a dialogue box. Risk profiles for variables other than operating margin can be substituted for the ones shown in the figure simply by double-clicking on the desired lines. While this paper is not about CFPM, brief descriptions of the model are provided in the sidebar.

The figure illustrates what is almost a universal feature of probabilistic models: the tendency for the variance of an uncertain quantity to increase over time. The operating margin appears on the *x*-axis and the frequency on the *y*-axis. All three graphs have the same scale. The means increase because of the positive growth rates assumed in CFPM, but none of the risk parameters change over time. The fifth percentiles—which depend on the mean, standard deviation, and any asymmetries in the probability distributions—provide a useful measure of downside risk. Getting used to looking at risk profiles and developing methods for mitigating risk should be an important goal for university planners.

CFPM reflects my conviction, based on decades of experience as a researcher, chief financial officer, and consultant, that a fully effective suite of analytical methods will be applied in colleges and universities only if the methods are easy to learn and convenient to use. Entering the data for the example required only a few minutes and the calculations only a few moments more. And while CFPM makes the process convenient and accessible to those without technical training, the Monte Carlo simulation methodology can be adopted with or without this platform.

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 expected profit determine a stock's price-earning 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 Comprehensive Financial Planning Model

The forecasting of revenues and expenditures is a key task for college and university financial planners. While the forecasts may be informed by data and other evidence, they invariably involve a great deal of judgment. Making the judgments in a coherent and consistent way can be difficult, especially when the variables are interconnected and the forecast must extend several years into the future. These difficulties are compounded when some of the financial outcomes are uncertain.

The challenge faced by planners is to visualize the consequences of alternative forecast and policy assumptions. This is easy enough when the number of variables is small, but more detailed models rapidly lead to overload. The complexity that comes with detail makes it difficult to manage changes in assumptions and understand the results. Ordinary spreadsheet models, the kind used by most planners, offer limited opportunities for visualization breakthroughs. Furthermore, most such models require formulas to be replicated many times (e.g., for each year in the planning period)—which increases workload and invites error. Worst of all, many complex spreadsheet models are opaque to all but technically sophisticated users, which limits policy makers' understanding and confidence.

Recent research on how to overcome these issues has produced the Comprehensive Financial Planning Model (CFPM) referred to in Massy's paper. The model

- Maintains a database of financial planning variables and displays several years of history for each variable.
- Embeds the university's "business rules," including risk management rules of the kind discussed in this monograph, within the financial plan to ensure consistency among planning variables.
- Optimizes the trade-offs among key planning variables while requiring convergence to "Long Run Financial Equilibrium" (cf. Hopkins and Massy, 1981, Chapter 6) at the end of the planning period if desired.
- Analyzes the consequences of financial risk and tests policies for mitigating risk.
- Saves and recalls tentative planning scenarios, propagates planning parameters to schools and/or departments, and rolls up results from schools and departments where applicable.
- Communicates results to diverse audiences using easily tailored report formats.

More information can be obtained from The Jackson Hole Higher Education Group.

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 non-profits 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.

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:

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

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. These restrictions can be relaxed, however, as will be discussed later.

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 it is easy to 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. A full-fledged capital structure model would include the size of the operating reserve, other reserves if applicable, and the details of the institution's debt covenants. We address these matters in the paper's final section. 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 paper.

I understand the difficulty of convincing trustee investment committees to change asset allocation targets, but such difficulties 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 (cf. Markowitz, 1952): 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

Table 1. Assumptions for the Operating Variables (thousands)

			Estimated Correlations				
Variables	*Mean*	*Std Dev*	*PvtEq*	*ForEq*	*DomEq*	*RealEst*	*Bonds*
Gifts for Operations	$20,000	$2,000	0.50	0.30	0.50	0.35	0.20
Incr. Research Revenue	$45,000	$3,375	0.10	0.25	0.30	0.30	0.10
Incr. Research Cost	$40,000	$2,000	0.20	0.10	0.20	0.15	0.30
The correlation between incremental research cost and revenue = 0.70							
Incr. Debt Principal	$10,000	na	na	na	na	na	na
Debt Service Rate							
Fixed Rate	2.5%	0.0%	0.00	0.00	0.00	0.00	0.00
Variable Rate: r=0	1.0%	2.5%	0.00	0.00	0.00	0.00	0.00
Variable Rate: r>0	1.0%	2.5%	0.00	0.00	0.20	0.40	0.60
Variable Rate: r<0	1.0%	2.5%	0.00	0.00	(0.20)	(0.40)	(0.60)

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 well-known book *Pioneering Portfolio Management* dismisses utility functions as impractical (Swensen, 2000, p. 126), but I'm convinced they can prove useful in the kind of analysis being conducted here. We shall use the following simple utility function:

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

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 current asset allocations are determined. 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 (cf. Swensen, 2000, pp. 126–127). I do not wish to get involved in arguments between financial theorists and practitioners but only note that preliminary tests of capital structure models can be conducted quickly and conveniently using the mean-variance utility model and refined later using simulation.

Tables 1 and 2 show the assumptions for the examples described in this paper. (All financial values are expressed in thousands of dollars.) The data are hypothetical, designed to illustrate the points being made. "Gifts for operations" is the subject of Example 1 and the other variables are added in Example 2. The variables will be described in their respective contexts. Values for the variables in Equation (1) that are not included in the tables are p = 5%, E_0 = $2 billion, and ONR = –$107.3 million in the base case. The ONR figure is the sum needed to bring expected margin to zero in the case being studied—a result that facilitates presentation. (See the last section for a discussion of how to compare fixed and variable rate debt.) In the "virtual endowment" case, for example,

Table 2. Assumptions for Asset Class Returns

				Estimated Correlations				
Asset Class		*Mean*	*Std Dev*	*PvtEq*	*ForEq*	*DomEq*	*RealEst*	*Bonds*
Private Equity	PvtEq	12.0%	28.1%	1.00	0.15	0.40	0.15	0.25
Foreign Equity	ForEq	7.2%	21.6%	0.15	1.00	0.50	0.16	0.25
Domestic Equity	DomEq	6.2%	15.1%	0.40	0.50	1.00	0.15	0.45
Real Estate	RealEst	3.7%	18.3%	0.15	0.16	0.15	1.00	0.25
Bonds	Bonds	2.1%	8.1%	0.25	0.25	0.45	0.25	1.00

ONR = –$107.3 – $20.0 = –$127.3 million.

Investment return depends on the parameters, standard deviations, and correlations in Table 2. I obtained the means and standard deviations by minimizing their squared deviations from the figures in David Swensen's Table 5.2, given the correlations (based on his Table 5.4) and the requirement that utility be maximized at the asset allocation he reports in Table 2.5 for large institutions (Swensen, 2000). The procedure also produced a value for ***k***, which turned out to be 5.04x10^4. Of course, these data are intended to be illustrative only.

Example 1: "Virtual Endowment"

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 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 2 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 3.

The first column of data in Table 3 shows our base

Table 3. 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%

case, in which there is no variation in the gift stream and, therefore, no correlations—which attempt to predict variations in gift streams—are used. 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 2's "Base Case" efficient frontier.

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.

In 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, loss of $479 (thousand) is more than compensated for in utility terms by the reduction of the standard deviation of operating margin by $705 (thousand), from $14,686 to $13,981.

Figure 2. Efficient Frontiers for Gifts to Operations, With and Without Correlations

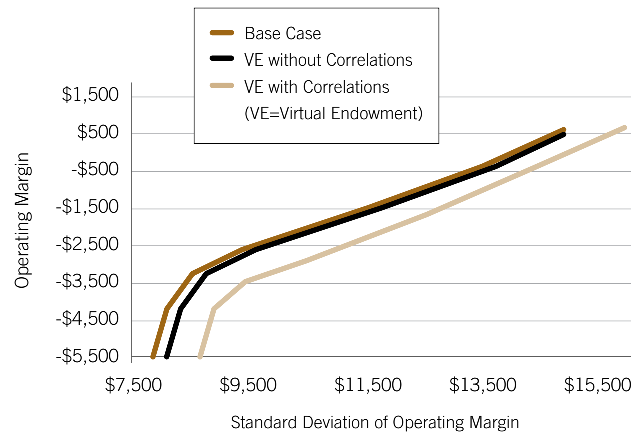

The example demonstrates clearly that an institution's optimal asset allocation can be quite sensitive to variations in operating revenue and expense—provided these variations are correlated with asset class returns. As for asset class returns, the revenue and expense variations will depend significantly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

Example 2: A Large and Risky Research Venture

Our second example deals with a large and risky research venture of the kind that might be undertaken as the result of a strategic planning exercise. The data were presented in Table 1. We assume that the venture is to be partly supported by gifts with the same risk profile as in Example 1 and that, in addition, there are incremental research revenues, costs, and debt service. Four alternative debt structures are presented: fixed-rate debt with a constant coupon rate of 2.5% and three different types of variable-rate debt, all with an expected payment of 1% of principle. I have structured the example so that expected incremental research revenue plus gifts equals expected incremental operating costs plus the coupon on fixed-rate debt.

The incremental costs and revenues of the new research program are uncertain, and of course so is the level of gifts. The correlations of research cost and revenue with the investment vehicles would depend on a combination of factors such as government appropriations, corporate budgets, and inflation. The less-than-complete correlation between the revenues and costs themselves reflects the fact that costs can only partially be adjusted in response to changes in revenues. Lack of experience with these correlations required the use of ad hoc judgments for purposes of the example. More detailed models of the relations among revenue, cost, and investments would sharpen the estimates—and likely generate some heated and potentially informative debates.

The fixed-rate debt alternative has a zero standard deviation and is uncorrelated with asset class returns because the coupon is assumed to be constant. The standard deviation for variable-rate debt was informed in a rough way by a statistical analysis of BMA Index data since 1980 supplied by Lehman Brothers. The data were presented as a log-normal distribution and the upside and downside one-sigma deviations are averaged. The applicability of this index is not clear, but it's likely that 2.5% is on the high end of the plausible range. No analysis of the correlations with investment returns has been done to my knowledge, but such a study surely would be both useful and practical. For the example, I simply chose some rather large positive and negative correlations to illustrate the effects.

While the example is hypothetical, it does have some important elements of reality. First, in contemplating such a research venture, most universities would insist that the program approach breakeven, including any expected incremental gifts—hence assuming breakeven simplifies the presentation without any loss of generality. Second, the debt alternatives are intended to span the range of possibilities planners may confront. They may choose between fixed- or variable-rate debt, and the latter may end up positively or negatively correlated with asset class returns depending on how the issue is structured. More realistic assumptions could be introduced without changing the essence of the example.

Table 4 shows the resulting asset allocations and statistics for operating margin in the same format as Table 3.

The first column of data in Table 4 shows the base case prior to undertaking the risky research venture. No correlations are used and asset allocations were determined by maximizing the utility function.

As in the earlier example, the following steps use the model to analyze the effects of various methods of funding the research venture:

Step 1 shows the effect of the risky research venture with no debt or fixed rate (FR) debt (which has zero variance as noted above).

Step1A, which holds the asset allocations constant, shows expected utility as having declined

Table 4. Effects of the Research Venture on Portfolio Statistics and Asset Allocations

			Step 1	Step 2	Step 3	Step 4
				Variable Rate Debt		
No Asset Class Adjustment	**Base Case**		**No Debt or FR Debt**	**R>0**	**R=0**	**R<0**
expected utility	100.0		75.5	76.1	75.5	75
expected operating margin	$0		$0	$0	$0	$0
std. dev. of operating margin	$13,545	Step 1A	$15,434	$15,399	$15,436	$15,474
expected total return	7.28%		7.28%	7.28%	7.28%	7.28%
std. dev. of expected total return	13.5%		13.5%	13.5%	13.5%	13.5%
With Asset Class Adjustment						
expected utility			77.3	77.9	77.4	76.8
expected operating margin			($724)	($750)	($724)	($698)
std. dev. of operating margin			$14,059	$13,969	$14,061	$14,152
expected total return			6.55%	6.53%	6.55%	6.58%
std. dev. of expected total return			12.3%	12.2%	12.3%	12.3%
		Step 1B				
Asset Class Allocations						
private equity	28.1%		26.8%	26.6%	26.8%	27.0%
foreign equity	21.1%		18.9%	18.6%	18.9%	19.1%
domestic equity	27.1%		17.3%	17.3%	17.3%	17.3%
real estate	13.3%		8.6%	8.8%	8.6%	8.3%
bonds	10.4%		28.4%	28.6%	28.4%	28.3%

due to the gift volatility and the revenue and operating cost volatility associated with the research. This volatility also shows up in the larger standard deviation of operating margin, which rose from $13,545 in the base case to $15,434.

Step 1B shows that maximizing utility with respect to the asset allocations can mitigate some of the operating risk: that is, the standard deviation of operating margin drops from $15,434 to $14,152. The asset allocations change because of the correlations between revenue and operating cost, on the one hand, and the asset class returns, and these changes serve to hedge operating margin.

Steps 2 through 4 add variable-rate debt service with three different assumptions about the correlation of the debt service with asset-class returns: Positive (R>0), zero (R=0), and negative (R<0). The results of the sub-steps, which are similar to the ones described above, can be summarized as follows.

The standard deviations and utility values change in the expected directions, i.e., negative correlations produce more risk, which lowers utility. The variable rate debt with positive correlations (R>0) structure produces the least overall risk because debt service enters the operating margin calculation with a negative sign, which makes the positive correlations produce offsetting effects. The effects are modest, however. The changes in asset allocation also are in the expected direction, but modest.

The conclusion that pops out from Table 4 is that the effect of the alternative variable-rate debt assumptions on the asset allocations is understandable but small. This is a consequence of my assumptions, of course, but the main ones are reasonably well grounded in reality. The main drivers are the amount of debt service and its standard deviation. Debt is already at half the level of endowment and, if anything, the standard deviation is on the high side at 2.5%—which still is small compared with the figures for most of the asset classes. One would have to increase the drivers dramatically to produce a material effect, even with the perhaps overly large assumed values of the correlations.

Several other conclusions follow from the example. First, the pattern of effects reinforces our finding from the virtual endowment case. Correlations between operating variables and asset class returns can generate significant opportunities for the hedging of operating risk through the endowment, but only if the size and volatility of the revenue and cost streams approach those of the investment vehicles. Second, the addition of fixed-rate debt does not affect asset allocation. (This runs counter to one of the conclusions from the Forum's 2006 Master Class, which was that the addition of debt of any type should be offset dollar for dollar in the endowment's bond portfolio.) Variable-rate debt can change the optimal asset allocation, but the effects are likely to be small.

Taken together, the two examples support the proposition that the effects of operational risk on optimal asset allocation should be quantified. This is true even if it may prove difficult to effect frequent changes in asset allocation. At the very least, the institution's base allocation should take into account the likely correlations between operating risk factors and investment return.

Downside Risk

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 (cf. Dorey, 2006; Hu, 2004; Natale, 2006).

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.

One important question concerns internal policies for debt capacity. To set ideas let me again go back to my time as Stanford's chief financial officer. The university was adding significant debt to finance new laboratory and other academic buildings as well as student housing. I believed we had plenty of cushion, a view that seemed to be shared by the bond rating agencies. Yet the board kept pressing my staff and me to "prove it." We had established a good reputation for modeling the university's financial structure and the trustees wondered, quite properly, whether we could extend the analysis to provide a rational basis for determining debt capacity.

The standard analyses focused on the ratio of debt principal to various asset measures—e.g., unrestricted quasi-endowment and other "available funds" items. We were in fine shape on all such measures, but both the trustees and I realized that it was the effect of debt service on the operating budget that was likely to be the limiting factor for our internal policy. It's worth noting in this connection that the rating agencies are concerned with payment of principal and interest, which for a well-endowed institution is likely to be less of a problem than the pain associated with operating budget adjustments. We ended up with a policy that set a numeric limit on debt service as a fraction of the operating budget, but the limit was determined by ad hoc judgment rather than a proper model. Now, these many years later, the tools are finally available to do the job right—for financially challenged institutions as well as for highly endowed ones.

Imposing a limit on the negative excursion of operating margin is straightforward, providing we can work with the margin's risk profile. The basic requirement is that margin does not go below the critical value more than a specified percentage of the time: i.e., that

$$\Pr[OM \leq OM_{crit}] \leq \Pr_{crit}, \qquad (3)$$

where "Pr" means probability. The equation is easy to apply when assumption (1) holds: Simply adjust the budget so that the expected value of operating margin minus 1.65 standard deviations equals the critical value. Departures from normality can be handled with Monte Carlo simulation. Again, one adjusts the budget until the desired percentile ($\Pr_{crit}$) of the calculated *OM* frequency distribution equals the limit.

Equation (3) works with any level of debt. In particular, it can be used to solve the debt capacity problem. One determines OM_{crit} in light of internal policies, rating agency preferences, and potential debt covenants and then runs the simulation for an ascending list of debt levels until the limit is reached. The analysis also can compare the use of incremental debt to build facilities with the appropriation of endowment capital gains or quasi-endowment principal for that purpose. If desired, the model could consider borrowing for investment in the endowment—a case that is important in for-profit theory. However, such actions may be precluded for endowed institutions by the anti-arbitrage rules for tax-exempt debt.

So far so good, but what about the possibility of mitigating operational risk by adjusting asset allocations? The problem is easy enough to formulate:

Maximize ***Utility*** with respect to the asset allocations while requiring that (4)

$$\Pr[OM \leq OM_{crit}] \leq \Pr_{crit.}$$

But while easy to formulate, the maximization cannot be accomplished with the procedure used earlier because no closed-form expression exists for $\Pr[OM \leq OM_{crit}]$. Approximations can be found, but a more direct way is to use Monte Carlo simulation, which has the added advantage of working when both of the assumptions given above are violated.

Combining the maximization in Equation (4) with Monte Carlo requires a technique called "chance-constrained programming" (a variant of stochastic programming; cf. Birge, 1997). In our case it works by generating a list of simulated operating margin values based on the asset allocations identified at each iteration of the maximization algorithm. Using the same random-number seed for all iterations makes the successive trials comparable and permits the algorithm to go forward. The fraction of cases in which *OM* falls below its threshold (i.e., the probability constraint) can be calculated directly from the list. The constraint is nonlinear, so a general nonlinear programming algorithm must be substituted for the more efficient quadratic programming procedure used in the examples. Such procedures now are readily available and quite easy to use. Indeed, the Premium Solver package offered by Frontline Systems, Inc. (www.solver.com), provides a powerful optimizer that can be seamlessly integrated with the company's efficient Monte Carlo system (Risk Solver) to obtain stochastic programming solutions in Microsoft Excel.

Table 5 shows some results from the SP model, for $\Pr_{crit} = 0.95$, using normal distributions for all of the variables and a sample size of 10,000. (Normality isn't necessary but it facilitates comparison with the results presented earlier.) The first column applies the SP model to our base case with no constraint applied. The results, which were computed in about half a minute, should equal those from quadratic programming as reported back in Table 3—and indeed they do, except for rounding errors. The second column applies a "light-touch" constraint: OM_{crit} = ($22,000) as opposed to ($22,306) in the unconstrained case. Scanning down the column we see that the mean, standard deviation, and expected utility have dropped slightly and that the asset allocation has tilted in the conservative direction.

More stringent constraints produce bigger effects, as shown in the remaining columns. For example, setting OM_{crit} at ($18,000) produces a strongly conservative asset allocation that drops utility from its original value of 100 to 65.4. The figures for "(fifth percentile – mean)/standard deviation (the last line of the table) show that the

Table 5. Downside Risk and Its Mitigation for the Base Case (thousands)

	No Constraint	*($22,000)*	*($20,000)*	*($19,000)*	*($18,000)*
Operating Margin					
5th Percentile	($22,306)	($22,000)	($20,000)	($19,000)	($18,000)
Mean	$4	($56)	($1,111)	($1,775)	($3,682)
95th Percentile	$22,235	$22,356	$18,132	$15,540	$10,596
Standard Deviation	$13,537	$13,503	$11,647	$10,636	$8,679
Utility	100.2	99.3	96.3	91.5	65.4
Optimized Asset Allocations					
Private Equity	28.1%	27.7%	21.5%	14.6%	7.6%
Foreign Equity	21.2%	20.2%	15.4%	14.5%	0.0%
Domestic Equity	26.9%	27.6%	25.0%	26.7%	18.9%
Real Estate	13.6%	13.8%	9.0%	8.5%	0.0%
Bonds	10.2%	10.7%	29.1%	35.6%	73.5%
(5th P'ctile-Mean)/StdDev	(1.65)	(1.63)	(1.72)	(1.79)	(2.07)

standard deviation drops faster than the constraint value as a result of the more conservative asset allocations. (The "1.63" in the second column is due to sampling or rounding errors.) Values of OM_{crit} much below ($18,000) cannot be obtained by adjusting asset allocation, so the stochastic program returns "infeasible" in these cases.

The Way Forward

The models 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 paper, but I will sketch the approach.

Use of multiple time periods will tend to amplify the above results because, as shown in Figure 1, uncertainty increases as one moves deeper into the planning period. Considering multiple periods would, among other things, reduce or eliminate the apparent dominance of variable-rate over fixed-rate debt. This will occur because the amplified variance will reduce utility and offset the lower interest rate. (The adjustments to *ONR* made in connection with Table 4 would have to be eliminated to get a true comparison between fixed- and variable-rate debt.) One would need to consider the serial correlations for the operating variables to get all of this right (investment returns generally are considered to be serially uncorrelated), but that would be a worthwhile task.

The multiperiod model can be adapted to substitute a minimum value for the institution's operating reserve in place of one for operating margin in Equation (4). The reserve balance at the end of a given year is simply the beginning value minus the sum of the margins through that year. One can maximize utility subject to a constraint on the reserve balance at the end of the planning period: zero, a positive value that provides a cushion, or a negative one that represents the limit of short-term borrowing. Other risk-mitigating procedures also can be incorporated—for example, the idea of a contingency or conditional budget like the one used at Stanford during my tenure as CFO (described in Hopkins and Massy, 1981, p. 282–284).

Mitigating the volatility of spending from endowment has been an important issue for university financial planners since the time spending was first linked to total return rather than the more predictable and stable yield component of return. Such mitigation usually is accomplished with a smoothing rule, such as a three- or five-year or exponentially weighted moving average, which shields the budget from short-term volatility. These models assume a predetermined asset allocation, but in principle the smoothing rule and asset allocations should be determined simultaneously as in the optimal control theory approach developed by Grinold et al. (1978). Except for CFPM, where smoothing was turned off to match the examples, the models considered so far do not include smoothing. However, it now may be possible to optimize the asset allocations and smoothing rule as part of a single model.

The next step in modeling might be as follows.

• Generalize Equation (1) to include "laws of motion" that link the planning variables in each period to those of the next. Represent the parameters of these laws as random variables in the Monte Carlo simulation. Equations based on growth rates will have the form $(1 + \rho_1)(1 + \rho_2)\ldots$, where ρ_t is the random growth rate for the *t*'th time period, but this is not a problem.

• Substitute log-normal distributions for the investment returns, and other variables as appropriate, in the Monte Carlo specification.

• Introduce serial correlations into the Monte Carlo as appropriate.

• Build the desired smoothing rule into the Monte Carlo.

Calculation of the variables in Equation (4) for a given Monte Carlo iteration will proceed without difficulty. Each of the variables is a function of asset allocation, which is assumed to remain constant through the planning period. The factors that multiply the allocation fractions will involve products of the growth rates, but for smoothing based on a moving average, the fractions themselves will enter the maximization, just as in the original problem. "Snake in the tunnel" and other discontinuous smoothing rules will change the structure. The maximization algorithm may work anyway, but if not one can try special algorithms designed to handle discontinuities. (Results not requiring the optimization of asset allocation will work with any smoothing rule.) All the analyses described in this paper, and many others as well, can be performed using the extended model.

I developed the stochastic programming procedure as a straightforward application of management science tools to solve the problem of mitigating downside risk. Later investigation showed a connection with deeper work by Boyle and Imai (2002), Huisman Koedijk, and Pownall (1999), Natale (2006), Sharpe (2007), and doubtless other experts on asset allocation—ranks in which this specialist on university management cannot be counted. The model does offer certain advantages of understandability and computability, however, and I hope it contributes to financial methodology in ways that go beyond the problems addressed in this paper.

References

Birge, John R. (1997). *Introduction to Stochastic Programming*. New York: Springer, Series in Operations Research and Financial Engineering.

Boyle, Phelim, Tan, Ken Seng, and Imai, Junichi (2002), "Asset allocation using quasi Monte-Carlo methods." Preliminary (March 4). http://en.scientificcommons.org/615300

Dorey, Martyn (2006). "Why aren't copulas more popular?" *Professional Investor* (February), pp. 18–21.

Grinold, Richard C., Hopkins, David S. P., and Massy, William F. (1978). "A model for long-range university budget planning under uncertainty." *Bell Journal of Economics* 9, 396–420.

Hopkins, David S. P., and Massy, William F. (1981). *Planning Models for Colleges and Universities.* Stanford, CA: Stanford University Press.

Hu, Ling (2004). "Dependence patterns across financial markets: A mixed copula approach." http://ideas.repec.org/a/taf/apfiec/v16y2006i10p717-729.html

Huisman, Ronald, Koedijk, Kees G., and Pownall, Rachael A. J. (1999). "Asset allocation in a value-at-risk framework." http://search.ssrn.com/sol3/papers.cfm?abstract_id=163970

Markowitz, H. (1952). "Portfolio selection." *Journal of Finance* 7, 77–91.

Massy, William F. (2008). "Cost and pricing in higher education." In *Handbook of Research in Educational Finance and Policy,* Edward Fiske and Helen Ladd, eds. New York: Routledge.

Massy, William F., Grinold, Richard C., Hopkins, David S. P., and Gerson, Alejandro (1981). "Optimal smoothing rules for university financial planning." *Operations Research* 29, 1121–1136.

Natale, Francesco Paolo (2006). "Optimization with tail-dependence and tail risk: A copula-based approach for strategic asset allocation." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=942275

Sharpe, William F. (2007). "Expected utility asset allocation." Working Paper, revised June 2007. www.stanford.edu/~wfsharpe/wp/euaarevised.pdf

Swensen, David F. (2000). *Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment*. New York: The Free Press.

William Massy is president of the Jackson Hole Higher Education Group. He is Professor Emeritus of Education and Business Administration at Stanford University, where he also served as chief financial officer. His most recent book is *Honoring the Trust: Quality and Cost Containment in Higher Education* (2003). Massy can be reached at bill@jhheg.com.