

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](http://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](mailto: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

costs are difficult to diversify, they induce a form of risk aversion in public companies, which induces these companies to take some risk management activities that reduce the volatility of firm value (the solid blue curve).

While many types of risk management tools exist, they can be grouped into three broad categories: (1) insurance, (2) asset–liability matching, and (3) swaps, futures, and derivatives. For-profits use all of these tools to mitigate risk, and intuitively the aim is to lower the volatility of net worth and therefore make it less likely that the firm violates a bond covenant. In general, all of these techniques are costly. If a firm buys insurance, its insurance premium covers not only the expected loss but also the insurance company’s overhead and profit. Similarly, there are transactions costs from executing swaps, and if one matches a floating interest rate liability with a floating-rate asset, the cost is the lower return on the floating-rate asset.

Massy’s key point is that the endowment potentially can be used as a fourth risk management tool.

Figure 2 shows that the operating balance sheet of a university looks like that of a for-profit, but in addition, the university has an endowment portfolio, which adds both assets and net worth. Note that the size of the endowment relative to the operating portfolio will vary across institutions, with some universities having very large endowments and others very small. This relative size will affect the costs and benefits to a given university when using the endowment as a risk management tool. For purposes of illustration, assume the endowment size as given in Figure 2.

When the endowment is used for risk management, it may be cheaper than other risk management tools that universities use. Suppose, for example, that a university is funding a major new building project. To fund the project, the university could (1) simply issue fixed-rate debt directly (to match the duration of the debt with the duration of the new buildings). But it might turn out to be cheaper to (2) issue floating-rate debt and then swap the floating interest payment to a fixed interest payment. Massy suggests another alternative: The university could (3) issue floating-rate debt but then offset that floating-rate exposure by adding floating-rate assets to its endowment. Which of these options the university would choose would be based on its assessment of which alternative would be least expensive and least risky. The key is that the university has added a heretofore unrecognized tool to its risk management toolbox.

### A Discussion of Risk Aversion

All risk management occurs because firms and institutions have reason to behave as if they are risk averse. Risk aversion occurs because people dislike losses more than they like gains. In the

Figure 1. Risk Management to Avoid Bankruptcy Costs

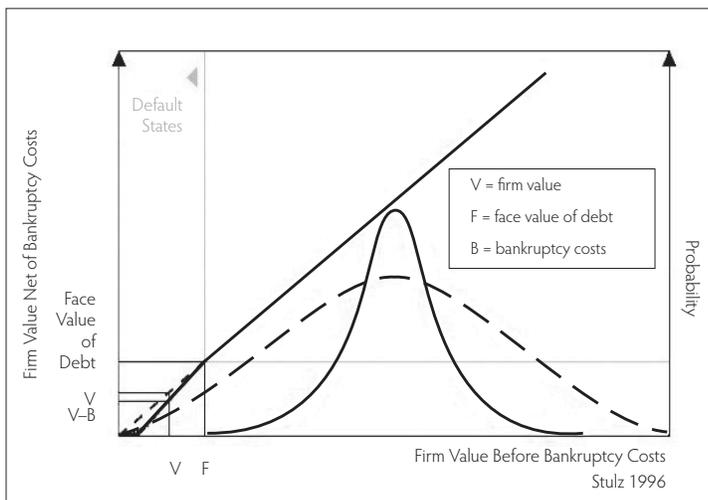
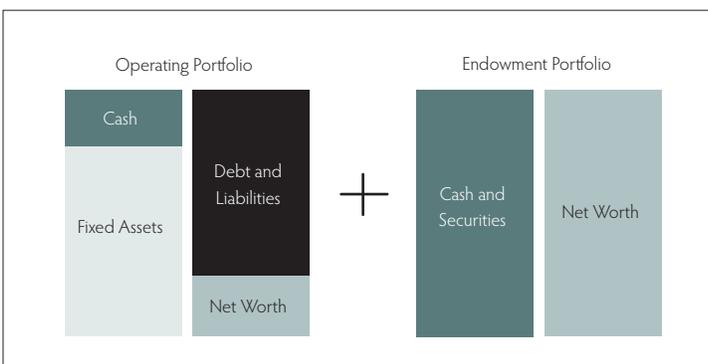


Figure 2. University Risk Management Tools



for-profit example above, risk aversion occurs because of bankruptcy costs: More bad things happen on the left side of Figure 1 when there is a very big loss (the bankruptcy costs) than do good things on the right side of Figure 1 when there is a very big gain. Universities as well are concerned about distress; having to sell assets to meet bond payments would be very costly, and most universities consider even a ratings downgrade a large cost to be avoided. Universities have an additional source of risk aversion: Unlike public for-profits, they have no analogue to public equity markets and thus are unable to raise money in public equity markets. Further, many universities strive to achieve “intergenerational equity” (Hansmann, 1990; Swensen, 2000) with respect to their endowments. When the university’s objective function involves student welfare, arguably the university’s risk aversion becomes that of the students. In this regard, students are typically fairly risk averse because they have no monetary wealth, only their risky future incomes. Similarly, one might argue that the university’s objective function should involve the interests of the faculty, the donors, and/or the trustees, and all of these individuals are likely to be fairly risk averse with respect to the university (particularly in comparison to public stockholders).

It is helpful to note that as a stockholder, I am not particularly risk averse because I can broadly diversify. If I read that the stock of the University of Phoenix, for example, is doing poorly, I am not greatly concerned because my mutual fund holdings give me small pieces of thousands of different companies. But if I care a lot about Yale, for example, there is no way for me to diversify my donation to Yale with donations to other charities or with a short position in Harvard. Anyone who cares about a charity cannot diversify that care and consequently will be risk averse with respect to that charity.

Massy uses the utility function  $U = E[\text{op. marg.}] - (1/k)SD^2$  as a straightforward way of capturing the idea that people like a higher return ( $E[\text{op. marg.}]$ ) but dislike higher risk, as shown by the term  $-(1/k)SD^2$ . The parameter  $k$  measures the individual's risk aversion. As  $k$  gets very large,  $(1/k)$  goes to zero, utility becomes unaffected by risk ( $SD$ ), and the person is risk neutral. On the other hand, when  $k$  is small,  $(1/k)$  is large, and the person dislikes risk, or is very risk averse. As Massy notes, to determine one's own  $k$ , one begins with one's current preferred asset allocation. Since this is preferred, utility should be highest for this asset allocation. To determine one's  $k$ , one should list a variety of different asset allocations, including the preferred allocation, and then fine-tune  $k$  so that the preferred asset allocation gives the highest utility,  $U = E[\text{op. marg.}] - (1/k)SD^2$ .

Choosing this utility function helps one order choices as one moves through alternatives. Figure 3 shows that when risk increases—as, for example, when an institution undertakes a risky research venture—the efficient frontier shifts to the right, from the gray line to the green line.

One choice is to just accept this extra risk and maintain the existing asset allocation, but this will not be optimal for a risk-averse university that wishes to reduce some of the increased risk. Depending on how risk averse the institution is, it will be

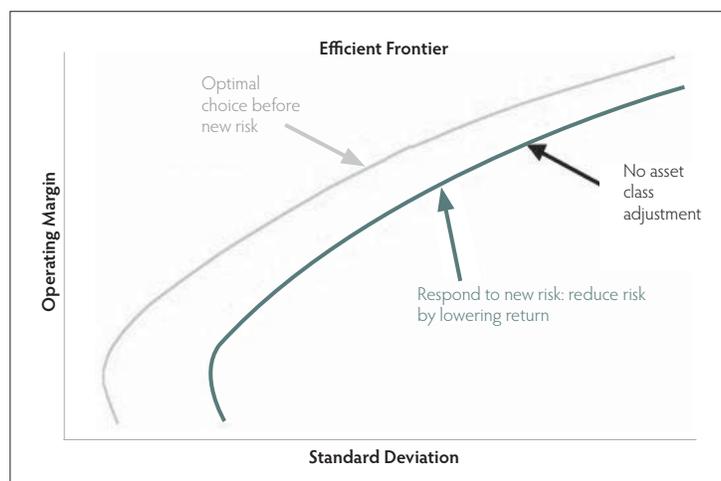
willing to give up some amount of return to lower risk. This is shown by the green arrow on the green efficient frontier, which is to the left (lower risk) and below (lower return) the initial status quo allocation. To reach that point, the university can use its endowment to reduce its risk (and increase its welfare or utility) by lowering its allocations in more risky equities and real estate and increasing its allocation in less risky bonds, as demonstrated by Massy's model.

### Relative Endowment Size and Risk Management

What would happen if the university's endowment were relatively much larger or smaller than the examples Massy uses? Organization risk aversion likely increases as the endowment becomes smaller relative to the university's operating budget. As a rule of thumb for comparing endowment sizes, create an endowment-to-expenses ratio by dividing the endowment by annual operating expenses. As an illustration, let's compare Yale and the University of Pennsylvania. At the end of fiscal year 2005, Yale's endowment was roughly \$18 billion. Penn's endowment, at \$6 billion, still seems fairly sizable until one realizes that Penn is a much larger operation and has annual expenses about twice those of Yale (about \$4 billion vs. \$2 billion). Thus, Yale's endowment-to-expenses ratio is about 9 to 1 and is about six times greater than Penn's endowment-to-expenses ratio of 1.5 to 1.

Because its endowment is so much larger, Yale likely behaves in a less risk-averse fashion than Penn. If Penn loses 20% or 30% of its endowment, it could very well bankrupt the entire institution. If Yale loses the same amount, its endowment would still be massive (at least by Penn's standards). So it would appear that Yale can take a lot more risk, both in its operations and its endowments.<sup>1</sup> So what Penn does, and what less well-endowed universities in general do, will likely reflect a greater degree of caution and risk aversion.

Figure 3. Effect of Increased Operating Risk



### Do Universities Already Do Something Like This?

The foregoing discussion suggests reasons why universities already may be implicitly using policies that amount to using the endowment as a hedging device. First, if a university has a

<sup>1</sup> While not exactly the subject of this discussion, the foregoing suggests a note of caution for readers of Swensen (2000). Because of its huge relative endowment, Yale can afford to undertake the high-risk, high-reward asset choices that Swensen does (because they can afford large short-term losses). Universities with endowments that are smaller relative to their operations may be too risk averse (or have too great consequences from short-term losses and illiquidity) to undertake these strategies.

small endowment, it is closer to its bankruptcy constraint than a university with a large endowment. To ensure that it has funds available if something goes wrong with its operations, it may invest its endowment more conservatively. Again, universities with large endowments can take risks and earn big returns; universities with small endowments may rationally take smaller risks and earn smaller returns. Second, universities with smaller endowments likely have larger relative virtual endowments. For example, Yale and Penn are in the midst of similarly sized \$3 billion-plus capital campaigns, suggesting their virtual endowments (future gifts) may be roughly the same size. Recall from above, though, that Yale's regular endowment is much larger than Penn's. Thus, Penn has a larger relative virtual endowment, which may also lead Penn to invest its endowment in less risky assets than Yale.

### Are There Reasons Not to Do This?

There may be reasons why an institution would want to maintain a separation of investment management and operations teams. One historical reason is the concern that too much of the endowment would be used for operations. Matters of governance may also come into play. Also, if the endowment is actively used to manage risk, this practice could possibly overly constrain the endowment portfolio managers. In this regard, Swenson cautions that there should not be too many people involved in asset allocation choices and that management by committee can spoil a portfolio manager's ability to earn returns.

### Conclusion

How might a university implement Massy's ideas going forward? Even if one is completely convinced by his ideas, it does not necessarily seem prudent to try to put them all in place in

the short term. Perhaps the best first step would be to just start thinking about correlations between operating budget items and asset returns. What is the correlation between future gifts to a university and various asset returns? At an institution such as Stanford, future gifts might be highly correlated with venture capital returns, whereas a university with a different alumni base might expect a correlation with a broader market return. If a university is research based, how will the magnitude of government and private grants vary with the economy and the stock market? Similarly, how much might the overall economy affect tuition revenue and costs such as faculty and staff salaries? It may be difficult to estimate these correlations, but Massy's point is that there is no reason to simply assume that they are zero. The danger of simply assuming that the correlations are zero is that a university may inadvertently be taking on a lot of risk—risk that could be avoided with slight adjustments to asset allocations in the endowment.

Massy's idea that endowments be used to manage risk is well worth adding to financial managers' toolboxes, and his model offers an excellent approach to implementing this important and new insight.

### References

- Hansmann, H. (1990). "Why do universities have endowments?" *Journal of Legal Studies* 19, 3–42.
- Stulz, R. (1996). "Rethinking risk management." *Journal of Applied Corporate Finance* (Fall), 8–24.
- Swensen, D. (2000). *Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment*. New York: Free Press.

John Core is the Ira A. Lipman Professor of Accounting at the Wharton School, University of Pennsylvania. Prior to joining Penn, he was an assistant professor of management at MIT. He was awarded the David Hauck Award for Excellence in Teaching (undergraduate division at Penn) in 2005. Core can be reached at [jcore@wharton.upenn.edu](mailto:jcore@wharton.upenn.edu).