Quantifying and Managing Endowment Portfolio Risk

One of the most important questions endowment managers and trustees have to answer is how much risk their endowment should assume through its asset allocation. We understand that this is not the simple question of deciding which asset categories are too risky, but the more sophisticated question of understanding how asset categories work in combination. An asset category that is high-risk on a stand-alone basis, may serve to reduce the risk of a portfolio when combined with other asset categories due to the nature of the correlations among the returns. Commodities are a good example. On a stand-alone basis, commodity returns have risks comparable to stock returns. However, since commodity returns have typically exhibited a negative correlation to stock returns, a portfolio built with a combination of the two exhibits less risk than either alone.

We begin by considering the range of available asset choices. We map the choices by plotting the risk and return characteristics of different asset classes and combinations of classes on a graph with return plotted on the vertical axis and risk plotted on the horizontal axis. We use the variability of returns as measured by standard deviation for the risk measure (see the risk primer box for why we use this measure). For the return measure, we use the average expected return from the class. The plots of some individual asset classes are displayed in Figure R.1A.[1]

A Map of Return and Risk Characteristics of Selected Individual Asset Classes

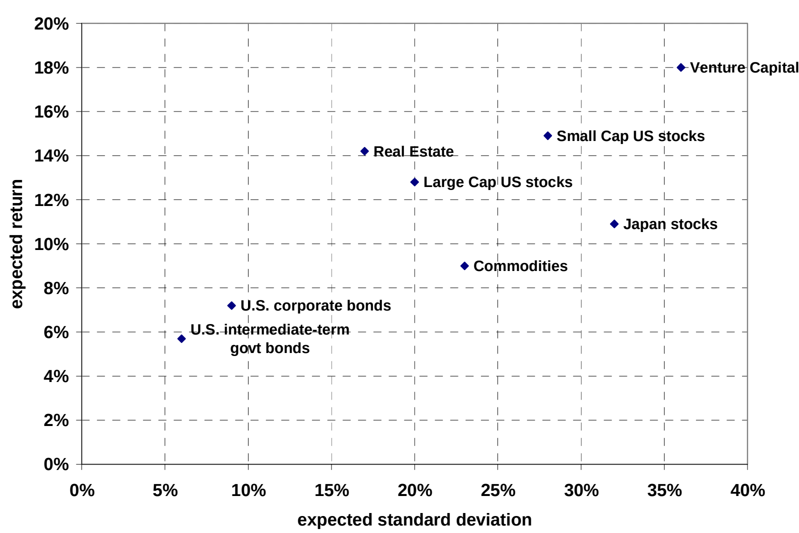

Figure R.1A: Return/risk Map of Selected Individual Asset Classes
Source: *Advisor*

Our primary interest is not in the choices among individual asset classes but in the choices among portfolios constructed by combining asset classes in various ways. For example, if we construct a portfolio divided 50/50 between U.S. large-cap stocks and intermediate-term U.S. Treasuries, we can expect the risk and return characteristics shown on Figure R.1B. Note that although the expected return on the mixed portfolio is just the weighted average of the expected returns of the two components, the risk is less than the average of the two because of how returns on the two classes have been correlated.

A Map of the Return and Risk Characteristics of a Blended Portfolio and Selected Individual Asset Classes

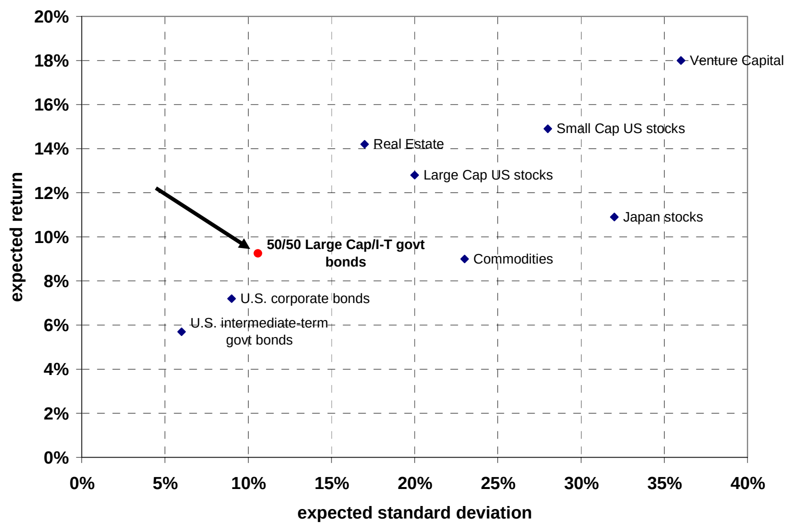

Figure R.1B: Return/risk Map of Selected Individual Asset Classes
Source: *Advisor*

If we add a third asset class to the mix, we get even more interesting results. If we add commodities, which have a lower expected return than our 50/50 portfolio and higher risk than either of the two components, we get a portfolio with a higher expected return and with <u>less</u> risk. Once again, this is due to how returns on the three asset classes have been correlated, with commodities showing a negative correlation to both stocks and bonds. The results are mapped on Figure R1.C.

A Map of the Return and Risk Characteristics of Selected Portfolios and Individual Asset Classes

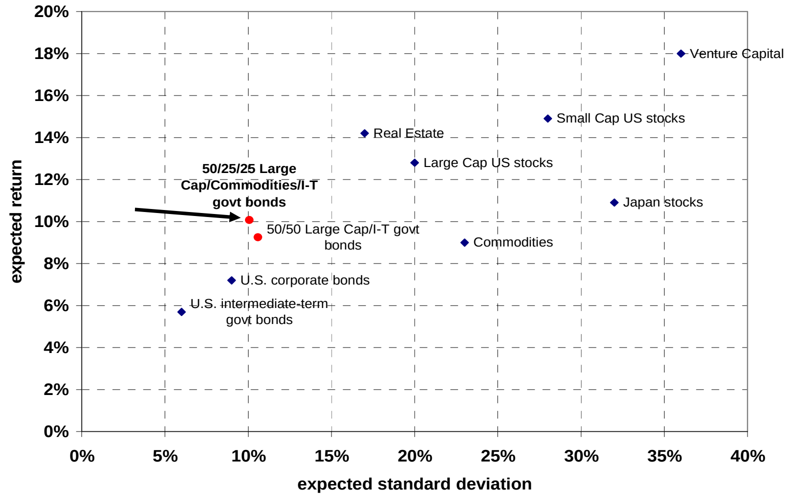

Figure R.1C: Return/risk Map of Selected Portfolios and Individual Asset Classes
Source: *Advisor*

As we add more asset classes to the mix and vary the amounts allocated to each, we get more interesting results. Figure R.2 shows the risk and return characteristics of more than 11,000 possible portfolios constructed from fifteen asset classes. The 11,451 portfolios mapped in Figure R.2 were constructed by varying the amounts allocated to each of fifteen asset classes in various combinations. Each point on Figure R.2 corresponds to the return and risk characteristics of a possible portfolio.

A Map of Return and Risk Characteristics of 11,451 Possible Portfolios Constructed from 15 Asset Classes

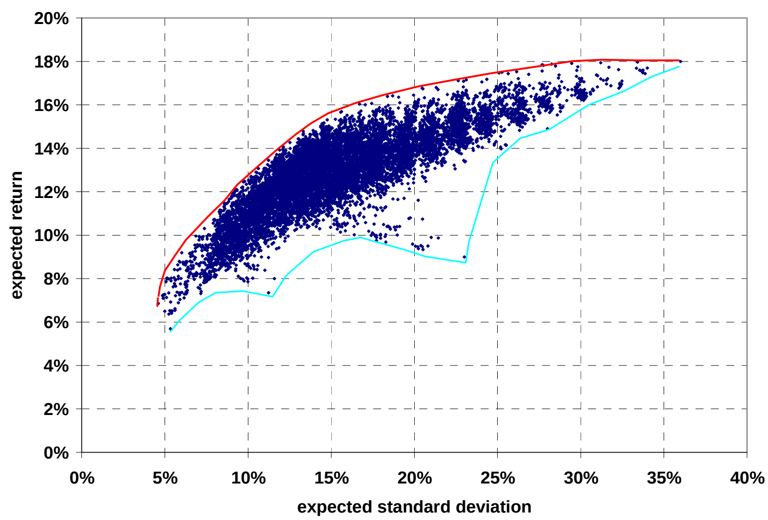

Figure R.2: Return and Risk Characteristics of 11,451 Portfolios Constructed from 15 Asset Classes

The possible risk and return outcomes are clearly bounded, as Figure R.2 shows. To the extreme right, outcomes are bounded by the point determined by a 100% allocation to the highest return/risk asset class, venture capital in this case. To the bottom, as indicated by the turquoise line on the chart, outcomes are bounded by those combinations that yield the lowest return for a given risk level—what we might term the least efficient frontier. To the top and left, as indicated by the red line, outcomes are bounded by those combinations that yield the highest return for a given level of risk—the efficient frontier. These are the optimal, most efficient combinations. For a given risk level, no greater return is possible from any combination of assets. Alternatively, we can look at the efficient frontier as mapping the lowest level of risk achievable for a given level of return. No matter how we combine assets, we cannot build a portfolio with less variability than that on the efficient frontier. The efficient frontier is exactly that; it is impossible to achieve a point above or to the left (to the northwest) of the efficient frontier line. It is impossible to construct a portfolio with higher returns or lower risk.

The question of how much risk we want to assume, then, comes down to the question of where along the efficient frontier our portfolio should be located. A low risk portfolio, one with a low degree of return variability, is located on the left of figure R.3. A high risk portfolio with a high degree of return variability is located on the right of figure R.3. If we want to eliminate all risk, we invest in Treasury bills and receive the rate of return at point A.

The Efficient Frontier, Return versus Risk

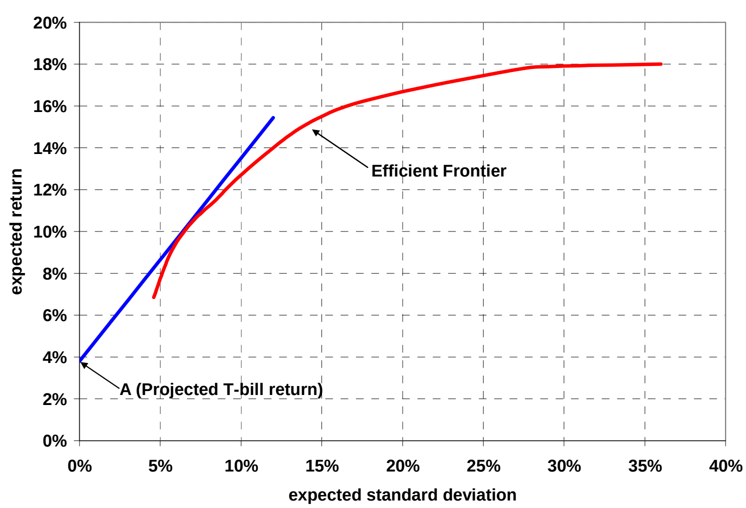

Figure R.3: The Efficient Frontier, Return versus Risk

But how do we decide where to be on the efficient frontier?

Consider first how not to make the decision, namely, by deciding on a target rate of return. It is tempting to set a desired rate of return and accept the risk that goes along with that return rate. This is tempting because most organizations have a better sense about their required rate of return than about their risk tolerances and exposures. An organization may know that to support their spending rate, cover inflation, and provide for some growth in endowment principal, they have to achieve at least an eight or nine percent return rate. If a ten or twelve percent return rate is achievable, all the better. Why not target the higher rate?

The problem with setting a target return rate is that the risk exposure accompanying that return may not be prudent. Deciding on an optimal asset allocation in this manner can lead an organization to take on excessive risk, to expose its operations to critical shortfalls in support, and to prevent it from coming to terms with and managing its real risk position.

To decide how much risk to assume or where to be on the efficient frontier, several facts about the organization are necessary besides knowing the minimum required return rate. The first is the amount of net discretionary spending in the organization's annual operating budget as a percent of annual budget support from the endowment. The second is how much of the endowment is comprised of original gift principal. These measures are the essential determinants of how much risk an endowment should assume. If they are unknown, organizations can only guess at their true risk tolerance. The previous two articles in this series explored the role of these two measures in determining risk. (See www.aametrics.com for downloads of the articles.) We shall see later that a third fact, average endowment giving is also relevant to an organization's risk tolerance.

Assuming we know the first two measures, how do they translate into a risk tolerance? Consider the example previously developed in the Part II of this series of an institution whose net discretionary spending (total discretionary spending net of that provided by endowment support) is equivalent to 6% of the total annual support provided by the endowment and whose endowment has accumulated earnings in excess of the original gift amounts of 15% of the total endowment value, or conversely, with an endowment 85% of which is comprised of original gift principal.

If we assume for a moment that budget support from the endowment is calculated on the basis of a spot endowment balance, a decline of more than 6% in the endowment balance would result in a drop in budget support of more than can be made up by reductions in discretionary spending. All other things being the same, either other sources of revenue would have to be increased, support for vital programs would have to be reduced, or liquid assets would have to be reduced.

A decline of more than 15% in the endowment balance would result in the loss of all budget support from the endowment because a 15% decline would reduce the endowment balance to the original gift principal and make withdrawals for spending support impossible.

Like all institutions, our sample institution has two risk thresholds to consider in endowment planning and risk management: the first determined by the nature of the institution's budget relative to endowment support and the second determined by the nature of the endowment itself. When the first threshold is crossed, endowment support is reduced to the point of requiring painful budget adjustments. When the second threshold is crossed, all endowment support is cut off.

The risk of crossing these thresholds depends on different factors. The risk of experiencing a decline in budget support depends on the variability of endowment returns, the nature of the spending formula used to calculate budget support, and the volume of regular endowment contributions. As we saw in Part II of this series, an average balance spending formula can smooth out endowment returns and provide some degree of protection from endowment market value declines. However, as we also saw in Part II, an average balance spending formula is not sufficient to protect an organization from all declines in spending support.

The risk of experiencing a cut-off of budget support depends largely on the variability of endowment returns. Although this is an issue where the law is evolving, the standard has been that original gift principal is inviolate. A market value decline that reduces the endowment principal value below the aggregate original gift principal cannot be smoothed out by average balancing. If the balance drops below the original gift principal, all draws on the endowment must cease until the endowment recovers. Although this may appear to be a remote issue for organizations with well established endowments, all organizations face this issue with newly established funds within an endowment.

Our sample organization, therefore, will seek an asset allocation which, in combination with its spending formula, minimizes the chances of a drop in budget support of more than 6%. And, it will seek an asset allocation which keeps the likelihood of a 15% drop in the endowment to an absolute minimum. The optimal asset allocation is one along the efficient frontier that has the highest expected return consistent with minimizing these risks.

Finding the optimal location along the efficient frontier is not a straight-forward task, but one that requires several steps. In technical parlance, it requires us to define our utility function and find where it meets the efficient frontier. We do this as follows.

First, we start with determining our minimum required rate of return. Ignoring risk for the moment, what rate of return do we need to keep even with inflation and to provide for the annual budget draw? For most organizations, this rate will be at least 9%. Although one might think that a lower rate of return would suffice for an organization with a 4% or 5% spending policy, the costs of most services provided by nonprofit organizations have increased at a rate higher than base inflation as measured by the consumer price index (CPI). Since human services do not lend themselves as readily to the productivity improvements that have benefited much of the rest of the economy, the cost of providing those services has risen more than average. The higher education cost index, for example, has increased at a compound rate of 5.2% from 1961 through 2007 during which time the CPI increased 4.3% per year.

We set the minimum return rate at 9.2% to cover a 4% annual spending allowance and a 5.2% annual increase in inflation. Organizations with higher spending allowances should set a higher minimum rate.

Nonprofit organizations, therefore, need to earn a minimum of 9.2% per year. Since this rate is well above the nominal risk-free rate—what we might earn on Treasury bills or guaranteed bank accounts—we have no choice but to turn to risky assets to achieve our target return. Once we turn to risky assets, however, returns become variable and our long-term compound return rate falls below our expected average rate of return (see the risk primer box on arithmetic and geometric rates of return). As returns become more variable, we need a higher average rate of return to be assured of a compound return rate that will meet our target. Figure R.4 shows the effect. When returns are modestly variable, say with a standard deviation of 10%, we require a portfolio with an expected average return rate of 9.7% to earn 9.2% on a compound basis. When returns become more variable, our expected return rate increases. At a 20% standard deviation, we need approximately 11.1%; and at a 25% standard deviation, we need over 12%.

Expected Return and Realized Compound Return as a Function of Risk

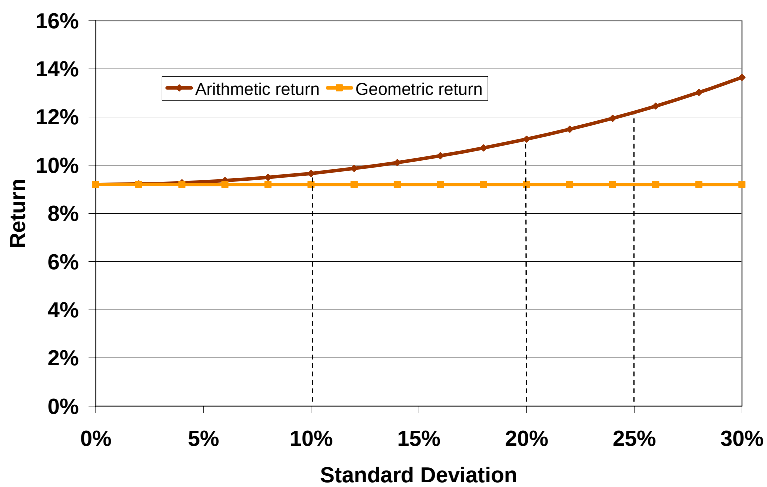

Figure R.4: Expected Return and Realized Compound Return as a Function of Risk
Source: *Advisor* calculations

Portfolios along the efficient frontier will be subject to the same return reduction due to the variability of returns. The greater the variability of portfolio returns, the greater the reduction in the compound rate of return. Figure R.5 shows how increasing variability reduces the compound rate of return of portfolios along the efficient frontier. Portfolios with the highest variability and the highest arithmetic returns actually have lower compound rates of return than portfolios with somewhat less variability. At a certain point, therefore, portfolios along the efficient frontier are no longer efficient from the perspective of maximizing the long-term compound rate of return.[2]

Efficient Frontier Expected Return and Realized Compound Return as a Function of Risk

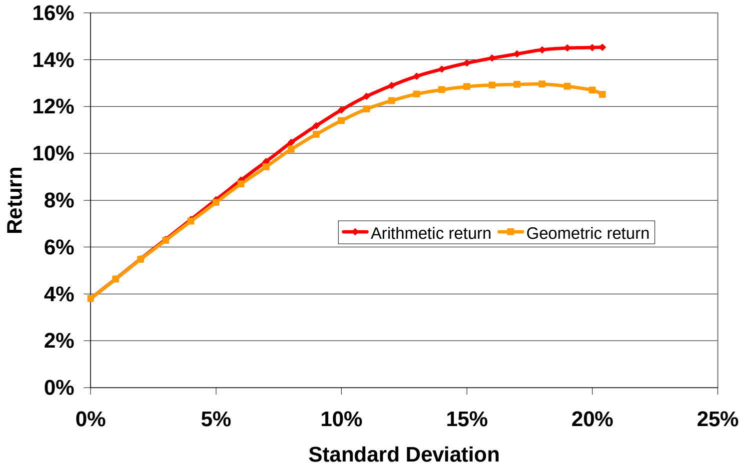

Figure R.5: Expected Return and Realized Compound Return from Portfolios along the Efficient Frontier as a Function of Risk
Source: *Advisor* calculations

The top line in Figure R.4 tracing the required returns to yield a 9.2% compound rate of return is our preliminary utility function—any portfolio with risk/return characteristics that fall above the line will meet our minimum required return. If we combine Figures R.4 and R.5, we can see how the risk/return characteristics of the portfolios along the efficient frontier compare with the risk/return characteristics of our minimum required return. Figure R.6 shows this comparison.

Comparison of Efficient Frontier Expected Return and Realized Compound Return with Minimum Required Return as a Function of Risk

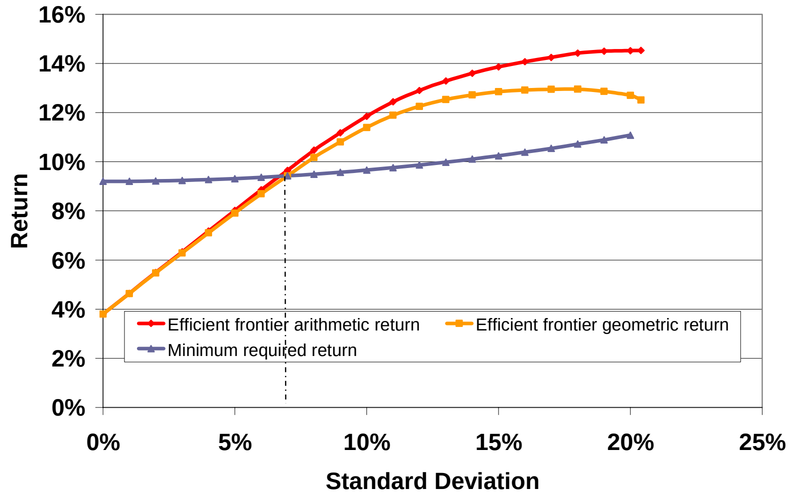

Figure R.6: Comparison of Efficient Frontier Expected Return and Realized Compound Return with Minimum Required Return as a Function of Risk
Source: *Advisor* calculations

As one can see from Figure R.6, efficient frontier portfolios with standard deviations of less than 6.7% do not generate a high enough return to meet our minimum requirement. All portfolios with higher standard deviations exceed the minimum return requirement, although portfolios with expected standard deviations of more than eighteen percent (18%) have declining geometric return rates. The optimal risk position for our sample case, then, appears to be in portfolios along the efficient frontier with expected standard deviations between 6.7% and 18.0%.

We can narrow the range further by factoring in the risk tolerances we specified earlier: minimizing the chances of a painful drop in budget support (of more than 6%) and minimizing the chances of a cut-off of budget support by minimizing the chances of a 15% decline in endowment balances. It is important to note that the best we can do is to minimize these risks. So long as we are invested in assets with variable returns, there will always be a risk of experiencing losses that will have painful consequences. We cannot eliminate the risks.

We cannot eliminate risks because variable returns not only reduce compound returns below average expected returns, they also introduce the possibility of loss. As results are more widely dispersed around the average, the possibility that some of those results are losses increases (see the accompanying box for a fuller explanation). Figure R.7 adds a curve showing the probability of falling below original gift principal at any year-end. As we would expect, the probability increases as the standard deviation relative to the expected return increases. The vertical dashed line shows the position of an efficient frontier portfolio with an expected return of 12.7% and a standard deviation of 11.5%. The probability of this portfolio experiencing a decline in year-end balance that would bring it below the original gift principal is approximately 2.5%, or once in every forty years. If we think this exposure is too high, we should adopt an efficient frontier portfolio allocation with a lower expected standard deviation.

It is well worth noting that the risk of a cut-off of endowment support due to a decline below original gift principal is a heavily time-dependent risk. The results reported in Figure R.7 are derived from the simulation of one thousand possible futures across the entire efficient frontier and are the highest probability for any length of time from one to two-hundred years. Since the likelihood of realizing positive returns in the first year is larger than the likelihood of experiencing negative returns, most simulated futures show an immediate increase in the accumulated earnings and market value gain, thereby reducing the risk of future performance reducing balances below original gift principal. Only those portfolios that experience initial losses have a meaningful risk of dropping below the value of original gift principal. On average, after the fourth year, the risk of dropping below the original gift principal declines.

Keeping in mind the time-dependent nature of this risk, if we can live with a one in forty chance of experiencing a cut-off of budget support due to a decline in balance below original gift principal, then we can set the upper limit on portfolio risk at the point along the efficient frontier with a 11.5% expected standard deviation. We have thus far narrowed the range of acceptable portfolios along the efficient frontier to those with expected standard deviations of between 6.7% and 11.5%.

Expected Returns, Realized Compound Returns, and Probability of Falling below Original Gift Principal (at 85% of beginning balance)

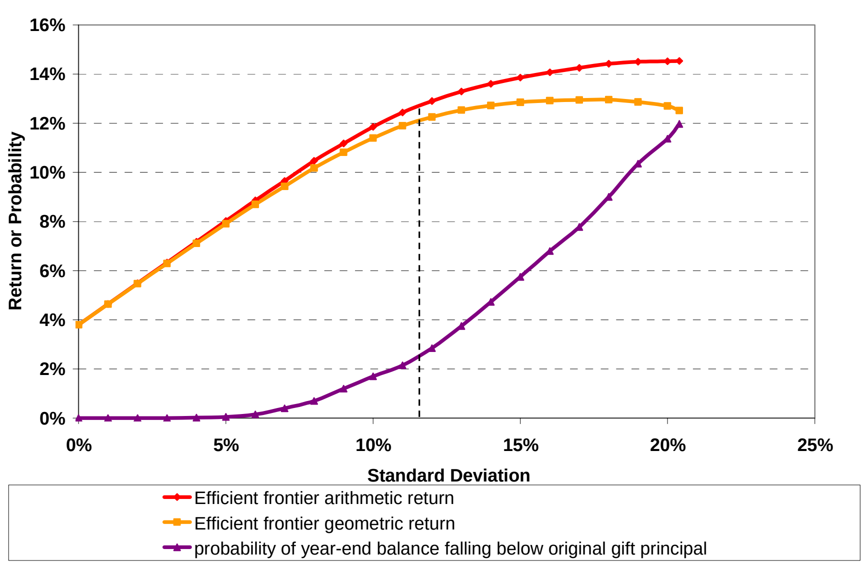

Figure R.7: Expected Return, Realized Compound Return, and Downside Return with 10% Probability as a Function of Risk
Source: *Advisor* calculations

The last element to factor in is the risk of suffering a spending drop larger than the amount of discretionary spending. In our example, this level is equivalent to 6% of the spending support currently produced by the endowment. A decline in budget support of more than 6% would require a reduction in vital programs, services, or support structure.

We can insulate ourselves from declines in budget support by using an average balance formula that smoothes out changes in endowment values by averaging balances buoyed by market value increases against balances depressed by market value losses. Although the typical average balance formula is effective at insulating endowment spending from the effects of temporary declines in market value, it is not sufficient to remove all risk, especially in periods of sustained market value losses. Figure R.8 shows the risk of experiencing a decrease in budget support of six percent or more in any year and assumes the use of an average balance formula that averages the previous three year-end balances to calculate the base for a four percent spending allowance.

Expected Returns, Realized Compound Returns, and Risk of a Decline in Annual Budget Support of more than Six Percent as a Function of the Variability of Returns

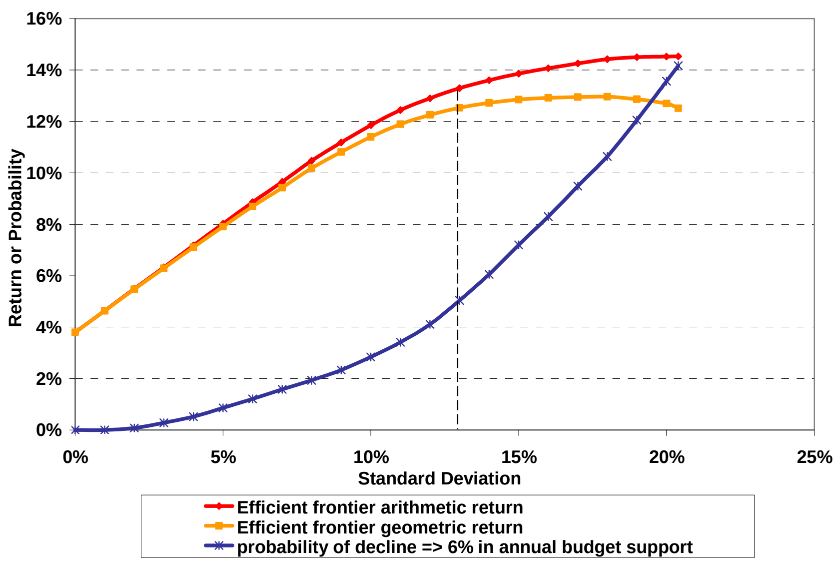

Figure R.8: Expected Return, Realized Compound Return, and Risk of a Decline in Annual Budget Support of more than Six Percent as a Function of the Variability of Returns
Source: *Advisor* calculations

The vertical line on Figure R.8 marks a portfolio with an expected return of 13.3% and an expected standard deviation of 13.0%. This portfolio has a 5% probability of experiencing a decline in annual budget support of six percent or more in any one year. In other words, it will experience budget drops of six percent or more on average in one of every twenty years.

The probabilities reported in Figure R.8 are exaggerated, however, because they do not take into account the risk reduction that occurs as a consequence of endowment contributions. Contributions ameliorate the risk of experiencing decreases in budget support.[3] Figure R.9 shows the probabilities of experiencing a decline in annual budget support of six percent or more in any one year under the assumption that annual endowment contributions are equal to one, two, or three percent of the prior year-end balance.

Expected Returns, Realized Compound Returns, and Risk of a Decline in Annual Budget Support of more than Six Percent at Various Contribution Rates as a Function of the Variability of Returns

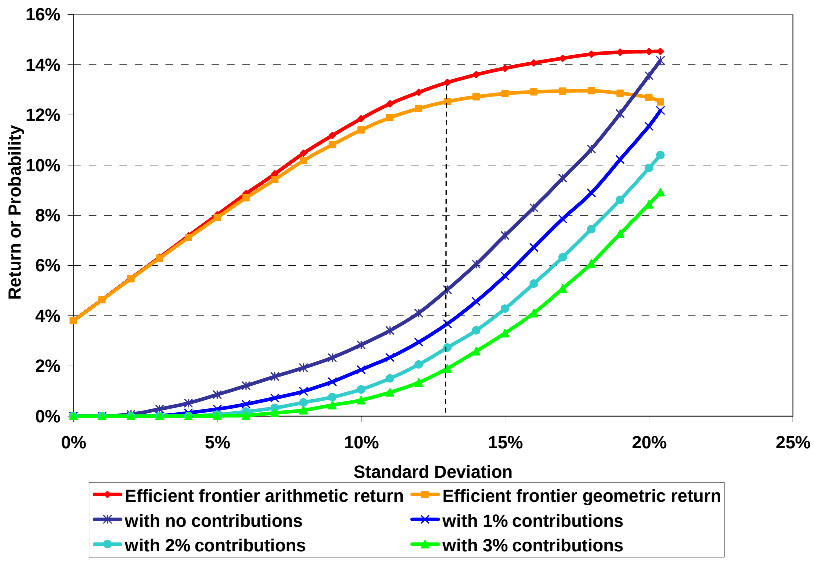

Figure R.9: Expected Return, Realized Compound Return, and Risk of a Decline in Annual Budget Support of more than Six Percent at Various Contribution Rates as a Function of the Variability of Returns
Source: *Advisor* calculations

As one would expect, for higher levels of annual endowment giving, the risk of experiencing a 6% or greater decline in budget support is reduced. For the portfolio we examined in Figure R.8, the probability of experiencing a budget decline greater than our 6% threshold drops from 5% to less than 4% at a contribution rate of 1% per year, to approximately 2.7% at a contribution rate of 2% per year, and to less than 2% at a contribution rate of 3% per year.

If we assume that annual endowment contributions in our sample case average 2% of endowment balance, we can combine all of the risk elements into one picture as shown in Figure R.10.

Comprehensive Portfolio Risk Assessment Showing Minimum and Maximum Acceptable Portfolio Risk

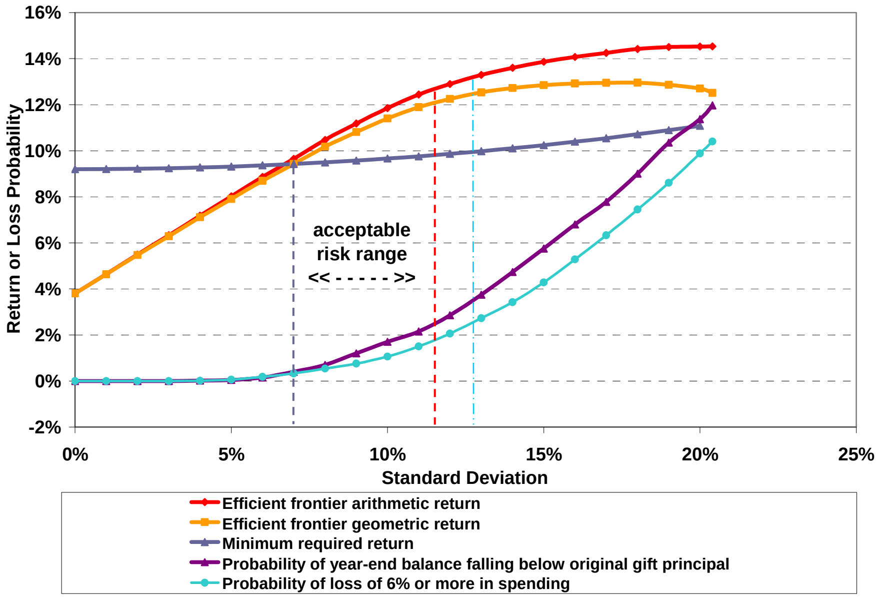

Figure R.10: Minimum and Maximum Portfolio Risk Limits for a Portfolio with a 9.2% Compound Return Requirement, and a Risk Tolerance of Violating Loss Thresholds for Annual Budget Support and Maintaining Original Gift Principal of No More than One-in-forty.
Source: *Advisor* calculations

Figure R.10 shows the limits on portfolio risk along the efficient frontier set by our three constraints: achieving a minimum rate of return, minimizing the chances of a painful reduction in endowment support of the operating budget, and minimizing the chances of a cut-off of all budget support due to a drop in endowment balance below original gift principal.

Achieving a minimum rate of return sets a minimum risk limit, indicated by the left-most, vertical, dashed line in Figure R.10. Any portfolio along the efficient frontier with a standard deviation of expected returns of less than 6.7% will not provide high enough average returns to sustain the purchasing power of the portfolio and provide for the annual support in accordance with the organization's spending policy.

Minimizing the risks of an excessive budget decline and a cut-off of all budget support sets the two upper risk limits, indicated by the right-most vertical lines in Figure R.10. In this case we have assumed a maximum risk tolerance for both of one-in-forty or a probability of 2.5%. In our example, the limit on maximum portfolio risk is set by the requirement to minimize the risk of a cut-off of endowment support since this limit (indicated by the red, vertical, dashed line) is less than the maximum risk allowed under our tolerance for excessive declines in annual budget support of more than six percent (indicated by right-most, vertical, dashed line). Any portfolio along the efficient frontier with an expected standard deviation of less than 11.5% will be within our upper risk limit.

The maximum portfolio risk under the requirement to minimize the risk of a decline in endowment support of more than six percent (6%) (as indicated by the right-most, vertical, dashed line) is not the limiting constraint in our case. In other cases, however, it may be. It might be if the accumulated earnings and market value gains constituted a larger proportion of the total endowment than in our case (15%) thereby reducing the chances that future losses would exceed the accumulated gains and earnings. It might also be if an organization relied more heavily on endowment support for vital operating expenses—in other words, if the amount of net discretionary operating expenses funded by the annual endowment draw constituted a smaller proportion of the annual endowment draw than in our case (6%).

The grey and red, vertical, dashed lines on Figure R.10 indicate the acceptable range for portfolio risk in our example, between 6.7% and 11.5% expected standard deviation. Where we choose to be within this range depends on an honest assessment of our risk tolerance and the trade-off between that risk tolerance and the portfolio return. If we can honestly live with one-in-forty odds of a cut-off of all budget support for the near term, then we should adopt the efficient frontier portfolio at the high-risk end of the acceptable risk range. Note that this portfolio (with an expected standard deviation of 11.5%) has an expected compound return rate of approximately 12.0% - nearly three percent more than our minimum requirement. Compounded over time, this excess return will result in a much larger endowment—nearly three times as large over forty years.

If one-in-forty odds are more than an organization can stomach, it can reduce the odds by adopting a lower-risk efficient frontier portfolio within the acceptable risk range. The portfolio with an expected standard deviation of ten percent (10%) reduces the risk of a cut-off of budget support to less than one-in-fifty (less than 2.0% probability) and the risk of a painful reduction in budget support to nearly one-in-ninety (1.1% probability). This portfolio still generates an excess return over the required 9.2% compound return, but with a smaller excess, approximately 2.2%.

We conclude with several observations about the process of measuring and managing endowment portfolio risk and asset allocation.

The first is that this case study demonstrates it is possible to assess an organization's portfolio risk in a way that is quantitative and meaningful to the management of endowment risk and asset allocation. Risk assessment does not have to be a vague and qualitative exercise. In fact, it should not be. Organizations have operating characteristics that determine their risk position in a strictly quantifiable manner. An organization can hardly go about setting risk tolerances and managing its endowment risk within those tolerances without going through the quantitative risk assessment demonstrated here.

The second observation is that just as it is possible to have too much portfolio risk, so is it possible to have too little. Maintaining the purchasing power of the endowment requires an average rate of return that can only be provided long term by investing in assets with variable returns, in other words, risky assets. Too little portfolio risk will not sustain the endowment's real value.

The third observation is that investing in risky assets inevitably means being exposed to the possibility of loss. We can limit our exposure to losses that would seriously harm the organization by keeping total portfolio risk within the acceptable range provided that we understand that the exposure to such losses cannot be eliminated. The upper risk limit set in Figure R.10 is based on a tolerance for experiencing painful losses no more than once every forty years. We can reduce that exposure by maintaining a portfolio risk exposure less than the limit of 11.5%, but we cannot eliminate the exposure, even if we adopt a portfolio with the minimum risk exposure.

Lastly, this example shows that generic risk limits for endowments don't exist. There is no one-size-fits-all analysis. Although the number of parameters that go into determining an institution's optimal risk position, or acceptable risk range, is not large, they can lead to significantly different results for different institutions. Prudent management of an organization's endowment leaves endowment managers and trustees no choice but to gather the facts and do the risk assessment and modeling necessary to determine quantitative risk limits.

---

[1] The data reported in Figure R.1A are annual average total returns and standard deviations derived from long-term total return indices for each asset class. The returns are average arithmetic returns.

[2] Astute readers will note that the efficient frontier displayed in Figure R.5 is not the same as in Figure R.3. In fact, both are derived from the same set of assumptions utilizing long-term arithmetic average rates of return, variabilities, and correlations. The efficient frontier in Figure R.5, however, imposes a minimum diversification requirement and does not permit all assets to be allocated to just one, two or three asset classes.

[3] Additional endowment contributions do not materially affect the probability of the endowment balance falling below the value of original gift principal, our first risk threshold. The incremental benefit of the additional earnings on the new contributions is not sufficient to offset the requirement to maintain the original gift value of the new contributions.