

Endowment Spending Policy An Economist's Perspective

Perry Mehrling
Barnard College

The majority of colleges and universities today target their endowment spending at about five percent of a 12-quarter average of total endowment funds. That five percent rate has been the standard since 1969, when the Ford Foundation recommended it in its well-known report, *Managing Educational Endowments*. Even though average endowment returns over the last 10 years were in the range of 10 percent, almost all endowment managers continued to hold to spending rules of about five percent. The result has been growing endowments, most often justified by the notion of intergenerational equity—that is, ensuring that future generations will be able to draw as much out of the endowment as we ourselves do today. Perry Mehrling, professor of economics at Barnard College, questions whether intergenerational equity requires endowments to *grow* by spending less than the income they generate, and proposes an alternative approach to endowment spending.

Intergenerational Equity: The Traditional Approach

Traditionally, the primary goal of endowment managers has been to maintain endowments in perpetuity to ensure a consistent and reliable level of investment income for generations to come—thus today's widely-used term, *intergenerational equity*. If we could be sure of the return on the endowment, it is clear what this conception of equity requires of us. We should spend the return and preserve the corpus. But the return is not certain.

Indeed the return behaves like a random variable, sometimes high, sometimes low, sometimes middling, and that fluctuation poses a problem for our concept of equity. If we want to guarantee to future generations a fixed amount of spending, then we cannot guarantee a fixed corpus. And if we want to guarantee a fixed corpus, then we cannot guarantee fixed spending. In a world of risk, it seems that we must choose between the putative interest of the beneficiary (constant spending) and the putative interest of the trustee (constant endowment).

Maybe not. One way to have it both ways (and so to avoid unpleasant boardroom squabble) is to fix the level of spending low enough that most of the time earnings are more than sufficient to cover it. In this way, we can guarantee both spending and corpus. Of course this means that spending will be modest, but on the bright side it also means that most years we will find ourselves with a bit of extra to do with as we will. In the course of events, we may wind up adding the extra to the endowment (so avoiding more unpleasant boardroom

squabbles). This is one way that endowments grow over time. (Another way is by new endowment giving, but we abstract from that here.)

This hypothetical policy is in fact the policy that most institutions have followed over time. The effect has been to ratchet up the corpus every time returns exceed spending, each time raising the level of assets that intergenerational equity requires us to preserve for future generations against any possible shortfall. The consequence can be measured approximately by the size of the so-called “quasi-endowment” relative to the “true endowment,” a number that today is in the range 25-30 percent for the average endowment.

***use the quasi-endowment slide (#8) here
rename it Figure 1***

Figure 1 shows that quasi-endowments, resulting from internal growth, comprise a sizeable portion of overall endowments and represent large sums of money.

The problem with the practice of adding excess returns to the endowment, as we are now finding, is that it only works if our spending rate is so low that excess returns are always positive. This is because we ratchet up the corpus when returns exceed spending, but we don't ratchet it back down when returns fall short of spending. Indeed, because every new higher level of achieved endowment accumulation becomes our new perpetual goal, the demands of intergenerational equity even seem to require us to make up any shortfall.

In other words, the policy of holding current spending below the expected rate of return is a policy that shifts all of the risk involved in future asset returns onto present shoulders, and none of it onto the shoulders of future generations. What is equitable about that? From a budget planning perspective, such a policy might well seem laudably conservative because it plans to spend only funds that are relatively sure to arise. Suffice it to say, though, that budgetary conservatism may be a different thing from intergenerational equity.

Intergenerational Equity: An Alternative Approach

There is another way. In the short run, at least, we can maintain both corpus and spending at constant levels by creating a *stabilization fund* to absorb fluctuations. Think of the fund as a kind of bank account with overdraft privileges, but one that has the same return as the endowment. On average the account is zero, but when there is a period of abnormally high returns, the account becomes positive,

and when there is a period of abnormally low returns, the account becomes negative.

Conceptually, we use the stabilization fund to take the risk of asset returns off the shoulders of both the present and future. With risk out of the picture, the requirement of intergenerational equity becomes clearer. We should set spending on the original corpus at the expected rate of return.

The remaining question is whether to adjust spending according to the size of the stabilization fund, and if so how. Clearly we cannot tolerate a stabilization fund that rises or falls without bound. The former would cheat the present and the latter would cheat the future. So we may want to spend from the stabilization fund in a way that stabilizes its value (at zero) over time. The main point however is that discussion about how best to treat the stabilization fund can be completely separated from the discussion about how to treat the original endowment.

Whatever approach is adopted, spending will be in accord with some version of the following basic formula:

$$S_t = \alpha E_0 + \beta F_{t-1},$$

where S_t is current spending, E_0 is the original endowment, F_{t-1} is the stabilization fund balance at the end of the previous fiscal period, and both α and β are spending rates.

To share investment risk equitably across generations, the spending rate α should be set at the expected rate of return. Of course, we don't know the expected return and so must estimate it, recognizing that the band of error is inevitably large. One response to that uncertainty is to underestimate the return "just to be safe." Hence, the traditional five percent rule. A major advantage of the stabilization fund approach is that it encourages more realistic estimates because the consequences of being wrong are mitigated by the dynamic operation of the spending rule. If we choose our estimate of α wrong, the effect over time will be a systematic deviation of the stabilization fund from zero, and then our choice of the adjustment coefficient β will come into play.

Comparison

To see how the dynamic stabilization approach might work in practice, consider an example in which the average return on a \$100 endowment is 10 percent, made up of equiprobable returns of 0 percent and 20 percent. Note that under this pattern of asset returns, the only way to ensure that the endowment never shrinks would be to fix the spending rate at 0 percent. The following table illustrates the dynamic stabilization approach under the assumption that we set

both spending rates, α and β , equal to the expected rate of return—in this example, 10 percent. Six periods are shown: three of 20 percent returns followed by three of zero percent returns.

Table 1: Dynamic Stabilization Approach, 10 Percent Spending Rate

<u>Endowment</u> (\$)	<u>Stabilization Fund</u> (\$)	<u>Income</u> (\$)	<u>Spending</u> (\$)
100	0	20	10
100	10	22	11
100	21	24.2	12.1
100	33.1	0	13.31
100	19.79	0	11.98
100	7.81	0	10.78
100	-2.97		

Compare this scenario with the traditional 5 percent payout, combined with a quasi-endowment ratchet when income exceeds spending, and borrowing for future replacement when spending exceeds endowment. In the traditional approach to intergenerational equity, with the same pattern of returns, we have the following scenario:

Table 2: Traditional Approach, 5 Percent Spending Rate

<u>Endowment</u> (\$)	<u>Quasi-Endowment</u> (\$)	<u>Income</u> (\$)	<u>Spending</u> (\$)	<u>Debt</u> (\$)
100	0	20	5	0
100	15	23	5.75	0
100	32.25	26.45	6.61	0
100	52.09	0	7.60	7.60
100	52.09	0	7.60	15.21
100	52.09	0	7.60	22.81
100	52.09			

To facilitate comparison of the two cases, Table 3 reports the totals at the end of the six periods.

Table 3: Summary Table

<u>Spending Rule</u>	<u>Total Spending</u> (\$)	<u>Terminal Fund Value</u> (\$)
10% + dyn. stabilization	69.17	97.03
5% + quasi-endowment	40.18	129.27

In the traditional case, endowment grows rapidly and as a consequence so does spending. Both endowment and spending reach a peak 56 percent above their

starting points, where they are maintained during the periods of low returns only by large borrowings. One can well imagine the state of mind of trustees and administrators at the end of the sixth period. Surely what is called for is a reduction in the spending rate, maybe from 5 percent to 4 percent, until the debt is repaid and the endowment restored.

The dynamic stabilization case, however, leads to a rather different conclusion. Here spending reaches a peak only 33 percent above its starting point, and the subsequent decline is buffered by drawing down the stabilization fund. At the end of the sixth period, the stabilization fund is more or less back to its long run zero target, and there is no apparent reason to change the projected spending rate of 10 percent. This is so even though more has been spent (\$29 more) and the endowment has grown by less (\$32 less) than the traditional case. Not only does the traditional rule systematically shortchange the present when returns are high but also, once it gets into trouble, the only apparent way out is to shortchange the present even more!

And there is one more thing in favor of the dynamic stabilization approach. Suppose we don't like the way the approach causes spending to fall when low returns eat into the stabilization fund. That doesn't mean we need to change α , the payout on the original endowment. We have set α equal to the expected return, and that is where it should be. If we are unhappy with the results, that means we need to change β , the payout on the stabilization fund. There is no particular reason that we need to set $\beta=10\%$. It doesn't even need to be a constant. Maybe we want β to be a function of the size of the fund itself.

Conclusion

The hypothetical traditional example above bears passing resemblance to actual experience in the last decade. In response to returns below spending, a large number of institutions have been compensating for negative investment returns by making sizeable fund transfers into the endowment from elsewhere in the budget. Others, fewer in number, are using new giving to make up some of the endowment shortfall.

No doubt immediate budgetary considerations are driving most of these responses, and probably that is as it should be. But the intergenerational equity approach may nevertheless provide a useful longer-term framework for understanding the current situation. The most important point is that equity does not require us to bear the burden alone. Quite the contrary, equity suggests that it is right and proper to get some help from future generations.

Concretely, it is useful to consider the quasi-endowment as a *de facto* stabilization fund that has ballooned all out of proportion to the original endowment. From this perspective, the problem in the last decade may not be that we spent too much, but that we spent too little! And the problem we face

now, given bad times for the last three years, is not how to preserve the quasi-endowment intact, but rather how optimally to spend it down.