A Robust Spending Rule: The α-β Approach

by Perry Mehrling

March 23, 2004

Introduction

The generally accepted perpetuity test of intergenerational equity proposes that we set spending from our endowments at a rate that can be sustained forever without invading the real value of the corpus. As a general rule, the test proposes that we set the long run rate of spending equal to the long run rate of return on our endowment. If we expect our endowment to generate 8% on average over the long term, then we fix the spending rate at 8%. The problem is that we don't know the long run rate of return. Whatever spending rate we pick is bound to be either too high or too low.

Given that it may be disruptive to an institution to set spending too high only to have to cut back later, we may prefer to set a rate that is from the start likely to be too low. If we reason in probability terms about all the possible future scenarios, instead of choosing a rate that will prove too high in 50% of the possible futures and too low in the other 50%, we may prefer to choose a rate that will prove too high in only 30% of the possible futures and too low in the other 70%. But that means giving up the goal of intergenerational equity since most likely we will be shortchanging the present. Is there another way of handling our uncertainty about the future?

One way would be to spend every period only the actual realized return in the previous period. If the return is 10% then we spend 10%, and if it is 5% then we spend 5%. By construction such a spending rule passes the perpetuity test. The problem is that the wide fluctuation makes rational budget planning very difficult. For this reason, the generally accepted test of budgetary soundness proposes that we set spending from our endowments at a <u>level</u> that can be sustained forever in real terms.

From a budgetary standpoint, as a general rule we want to set the level of real spending equal to the long term real dollar return. If we expect our $100 million endowment to generate $8 million on average over the long term, after taking account of inflation, then we fix the level of spending at $8 million and increase it every year only by the rate of inflation. The problem is that we don't know the long run real dollar return. Whatever spending level we pick is bound to be either too high or too low, and if we try to protect the future by choosing a spending level that is likely too low, then we wind up shortchanging the present. Once again, is there another way?

The great weakness of these usual approaches to formulating endowment spending rules is that they require us to pretend that we know more about the future than it is possible to know. We do not know the long run return, and we cannot estimate it with any acceptable degree of accuracy. What we need is a rule that is robust in the sense that it works reasonably well even if our guess about the long run return turns out to be wrong. The α-β approach provides just such a robust rule.

The α-β Approach

The key to the α-β approach is to separate the total endowment into two distinct funds. The investment of these two funds can be the same, so the separation is mainly a matter of accounting. The first fund is the original endowment, denoted E, which we allow to grow at the rate of inflation. The second fund is the difference between the total endowment and the first fund. We call this second fund the stabilization fund, denoted F. By construction, the original endowment fund remains constant in real terms over time, and all fluctuation in endowment value is absorbed by the stabilization fund.

Once we have the two funds, we apply the spending rate α to the first fund and the spending rate β to the second fund to get our spending rule:

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

Note that the two spending rates can be different.

The perpetuity test of intergenerational equity says that we should set α equal to the expected return, so that is the same as the traditional spending rules. Suppose for concreteness that we think the expected return is 8%, but are only comfortable committing to spend at a rate α=6%. Then, if we are right about the 8%, we will wind up spending too little and our mistake will show up as a growing stabilization fund.

Turning to the stabilization fund, the perpetuity test of intergenerational equity says that we should keep it near zero in the long run, and that's where the second spending rate comes in. In order to keep the stabilization fund near zero, we need to set β greater than the true expected rate of return. If we think the expected return is 8%, we might want to set the spending rate on the stabilization fund at 10%. Then if we are right about the 8%, we will be spending out of the stabilization fund at a rate higher than the natural growth of the fund. Since the stabilization fund will tend to shrink over time, we pass the perpetuity test of intergenerational equity.

But what if we are wrong about the 8%? Suppose the true long run return is only 6%. In this case, the spending out of our original endowment is just right, so the stabilization fund remains at zero, and once again we pass the perpetuity test. Suppose the true long run return is only 4%. In this case, the spending out of our original endowment is too high, so the stabilization fund tends to become negative. Spending out of the stabilization fund is therefore also negative, and this works to correct the overspending out of the original endowment. Again, because the spending rate is higher than the true return, the stabilization fund will tend toward zero over time, and once again we pass the perpetuity test.

The robustness of the α-β rule can be seen more generally in Figure 1, which uses historical returns taken from the NACUBO Survey to show the consequences for the stabilization fund of different parameter choices for α and β. The three lines at the top show cases where α=5%, which was an underestimate of the actual return over the period. The three lines at the bottom show cases where α=9%, an overestimate. The lines in the middle show cases where α=7%, which was just about right. In each case, it is clear how high values of the β parameter help to correct mistakes in the α parameter.

Regardless of the choice of alpha, the line corresponding to the highest beta value keeps us closest to the target value of zero.

Figure 1

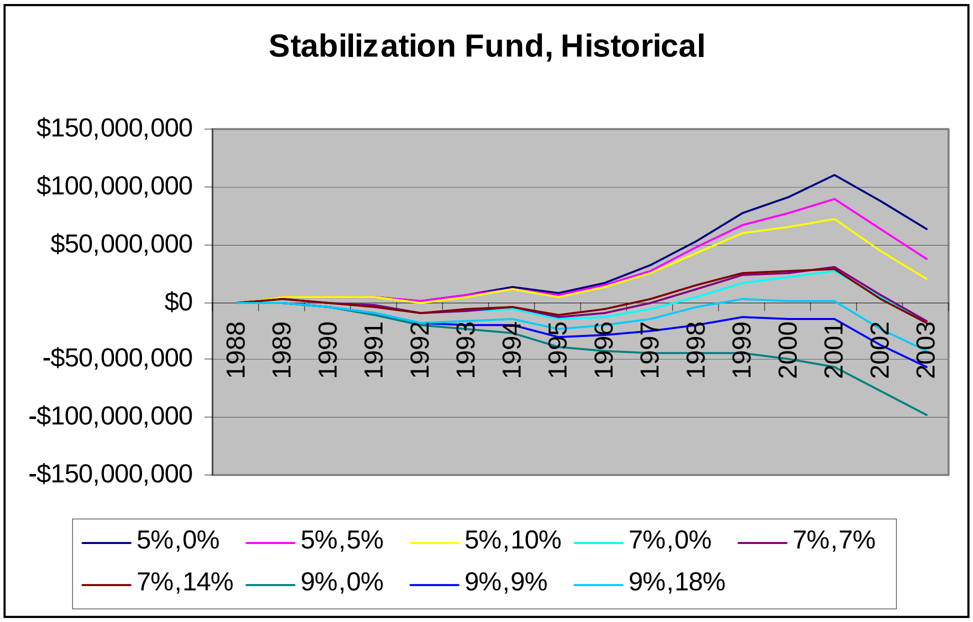

Figure 1 also enables us to compare the α-β rule with traditional rules. The rule to spend a constant fraction of the total endowment is essentially an α-β rule with $\alpha=\beta$, for example $(\alpha,\beta)=(5\%,5\%)$. The rule to spend a constant total amount is essentially an α-β rule with $\beta=0$, for example $(\alpha,\beta)=(5\%,0\%)$. Both of the traditional rules perform fairly well so long as we get α about right, but are clearly inferior if we get α wrong.[1]

[1] More generally, if returns are independently distributed over time, then intergenerational equity requires $\beta>\alpha$ even if we get α exactly right.

Table 1 reports on the robustness of the α-β rule looking forward, using the Commonfund Allocation Planning Model to generate 1000 different scenarios for the next 20 years. The portfolio allocation is the average for endowments of $1 billion and or more, according to the Commonfund Benchmark Survey. For this portfolio, the assumed median annual real return rises over twenty years from about 5% to almost 8%, with an average about 7%. We focus on the ending value of the stabilization fund measured as a proportion of the ending value of the original endowment. For all choices of α, higher values of β bring the median stabilization fund closer to zero, and reduce the variance around that median as well.

Table 1

Alpha	Beta	Real Spending			Real Fund Value		
		1st Q	**Median**	**3rd Q**	**1st Q**	**Median**	**3rd Q**
4%	0%	4%	4.00%	4.00%	12.10%	105.30%	246.60%
	4	4.30%	5.30%	7.00%	8.30%	68.10%	152.90%
	8	4.50%	6.10%	8.50%	4.10%	42.90%	98.20%
5%	0%	5%	5.00%	5.00%	-21.10%	60.10%	194.40%
	5	4.80%	6.00%	7.80%	-12.10%	36.30%	105.20%
	10	4.80%	6.50%	9.00%	-9.60%	19.90%	61.80%
6%	0%	6%	6.00%	6.00%	-53.40%	19.60%	141.60%
	6	5.20%	6.50%	8.40%	-28.90%	10.30%	66.00%
	12	5.00%	6.70%	9.20%	-18.90%	3.60%	35.00%
7%	0%	7%	7.00%	7.00%	-87.80%	-23.80%	93.30%
	7	5.60%	6.80%	8.80%	-42.60%	-10.90%	34.00%
	14	5.20%	6.80%	9.20%	-26.10%	-7.70%	15.90%

The α-β rule is apparently a robust method for achieving long run intergenerational equity. Figure 2, which shows the path of spending for the same sequence of returns that produced the middle lines in Figure 1, illustrates the major drawback of the rule. The α-β rule tends to produce fluctuation in spending that will make a treasurer's life difficult. Can something be done about that without loss of intergenerational equity?

Figure 2

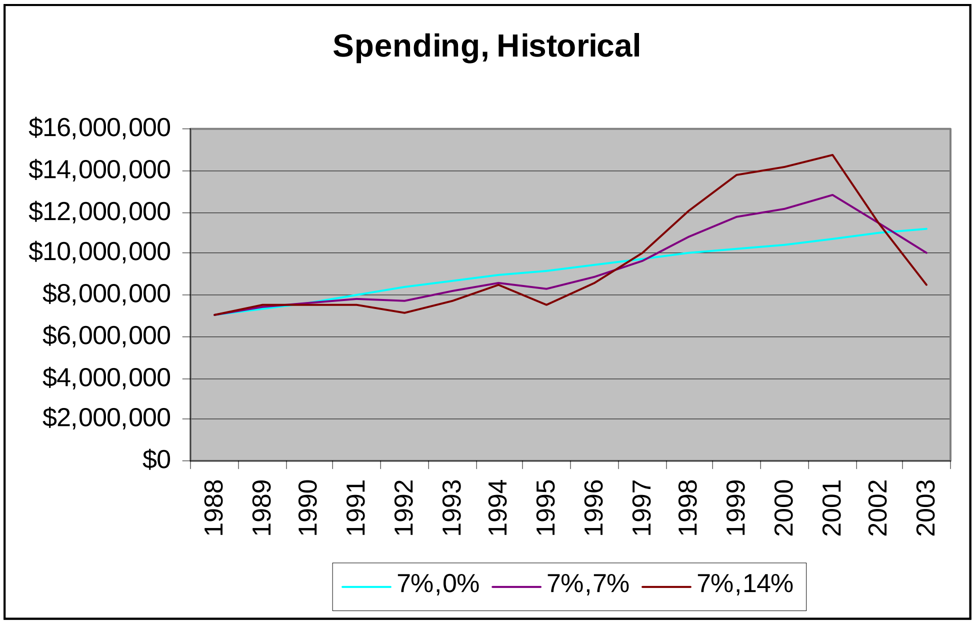

Smoothing Rules

It is important to be clear about what exactly we are trying to smooth. Table 2 reports the variation of cumulative real spending, where spending is measured as a proportion of the original endowment. Clearly higher values of β cause greater variation in realized cumulative spending across the different scenarios. But arguably that is just as it should be.

We want spending to track the long run realized rate of return. In the Commonfund projections, the cumulative return has a fairly wide variance, ranging from a first quartile value of 4.59% to a third quartile value of 9.11%, around a median of 6.91%. What we want is a rule that causes cumulative spending to exhibit a similar pattern. From this point of view, clearly higher values of β are better. Not only do they keep the stabilization fund near zero, but also they keep spending in line with long run return.

To gain intuition about how this works, Figure 3 plots cumulative spending against cumulative return for each of the 1000 scenarios following the spending rule α=5%, β=10%. Apparently this α-β rule produces fairly good long run correlation between spending and return, though there is some indication of overshooting for extreme return realizations. That long run correlation is exactly what we want, and we don't want smoothing to eliminate it.

Figure 3

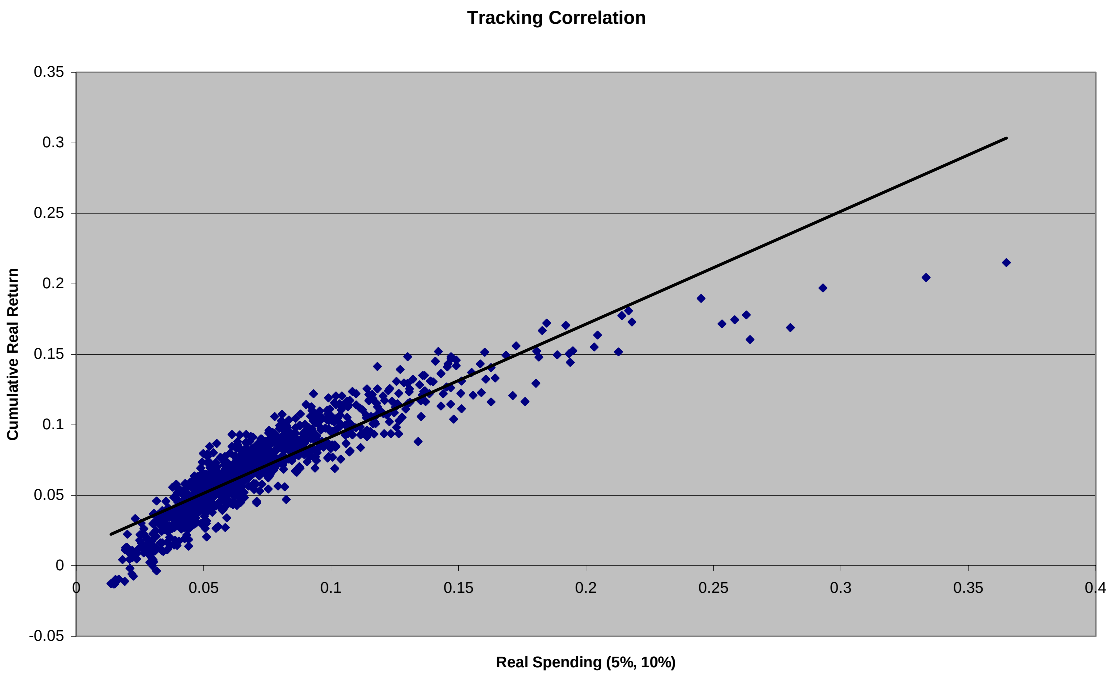

What we do want is to smooth the year-on-year fluctuation of spending. Two kinds of smoothing are already in wide use in conjunction with traditional spending policies. Both kinds can also be implemented in conjunction with an α-β rule.

Most institutions apply their chosen spending rate to a three year, or twelve quarter, average of their endowment's total market value. Obviously there is no point to such a procedure for the original endowment, but we could easily use it to smooth the spending from the stabilization fund:

$$S_t = \alpha E_{t-1} + \beta[F_{t-1} + F_{t-2} + F_{t-3}]/3.$$

A more elaborate form of smoothing is embodied in the Stanford Rule, a version of which is used by many other university endowments. The Stanford Rule sets spending at a weighted average λ of the previous year's level of spending and (1-λ) of the target spending rate times the total endowment market value. The effect of the rule is to apply the target spending rate to a weighted average of endowment market values stretching into the distant past, with more recent values weighted more heavily. Such a smoothing scheme could also be used with the α-β rule as follows:

$$S_t = \lambda S_{t-1} + (1-\lambda)[\alpha E_{t-1} + \beta F_{t-1}].$$

Refinements

So far we have treated β as a constant, and emphasized the importance of choosing β greater than the expected return. More sophisticated formulae may be worth considering. Table 2 reports results for the following three cases, again using the Commonfund APM to project performance for the next 20 years. In all three cases, we set α=5%.

Case 1: Asymmetric β. Here we apply a spending rate to the stabilization fund only if the fund is positive.

$\beta = 0, \quad F_{t-1} < 0,$
$\beta = 10\%, \quad F_{t-1} > 0.$

Case 2: V-shaped β. Here we increase the spending rate depending on the size of the fund's deviation from zero.

$\beta = |.2*F_{t-1}|$

Case 3: Banded Inflation. Here we keep the overall level of spending constant until the overall spending rate hits either 4% or 6%.

$\beta = 4\% - (\alpha-.04)E_{t-1}/F_{t-1}, \quad .04 = S_t /(E_{t-1} + F_{t-1}),$
$\beta = 0, \quad .04 < S_t /(E_{t-1} + F_{t-1}) < .06,$
$\beta = 6\% - (\alpha-.06)E_{t-1}/F_{t-1}, \quad .06 = S_t /(E_{t-1} + F_{t-1}).$

Table 2

Alpha	Beta	Real Spending			Real Fund Value		
		1st Q	Median	3rd Q	1st Q	Median	3rd Q
5%	10%	4.80%	6.50%	9.00%	-9.60%	19.90%	61.80%
	(0, 10%)	5.20%	6.50%	8.90%	-30.10%	14.00%	56.70%
	|.2*FND|	4.90%	6.40%	9.60%	-12.70%	22.70%	49.80%
	Banded	4.90%	5.40%	6.90%	-15.00%	51.90%	138.50%

The simulations show that, compared with the case of β=10%, the asymmetric beta case produces a more symmetric distribution of fund outcomes. The V-shaped beta case produces a somewhat tighter distribution. The Banded Inflation case does worse than the others because beta is always less than the long run rate of return.

Conclusions

Traditional spending rules that fix a constant spending rate or a constant spending level do not perform well when the realized return deviates from the expected return. The α-β approach makes clear the source of the problem. The implied spending rate on the implied stabilization fund is too small.

Implementation of the α-β approach is straightforward. First we calculate the implicit stabilization fund that has already been generated by whatever version of one of the traditional rules we may be following. That calculation will produce a historical record like one of the lines in Figure 1.

Second, we choose the spending rate on the stabilization fund deliberately with an eye to intergenerational equity in the face of uncertainty. We can choose β to be a constant, or a function of the size of the fund. Here we can be guided by the kind of simulation results reported in Tables 1 and 2.

Third, we apply one of the traditional smoothing rules. Since higher values of β tend to produce greater fluctuation in spending, this smoothing step is especially important. Three-year averaging may not be enough and some version of the Stanford Rule may be desirable.