How to Squander Your Endowment: Pitfalls and Remedies*

Philip H. Dybvig[†]

Washington University in Saint Louis

Southwest University of Finance and Economics

Zhenjiang Qin[‡]

University of Macau

October 8, 2021

*We are grateful for comments from Michael Brennan, John Campbell, Thomas Gilbert, Hong Liu, Jun Liu, Mark Loewenstein, Ian Martin, and Xudong Zeng, and for seminar comments at the AFA, CICF, EIEF, Fordham, University of Macau, University of New Orleans, Renming University, SAIF, SJTU, SWUFE, Tsinghua, Ulm, Wash U, and Wisconsin. We are grateful for support from the EIEF and SWUFE, and for support from the University of Macau under grants SRG2018-00113-FBA and MYRG2018-00210-FBA.

[†]Address: Washington University Olin School, Saint Louis, MO 63130-4899, USA. Email: dybvig@wustl.edu

[‡]Corresponding author. Address: E22, Faculty of Business Administration, University of Macau, Taipa, Macau, China. Email: zhenjiangqin@um.edu.mo

Abstract

University donors can contribute to endowment to make a permanent impact, given that universities are committed to investment and spending policies that preserve capital. Unfortunately, practitioner's common criterion of spending less than the expected return on endowment is known not to be sufficient to preserve capital, and we show that smoothing spending is destabilizing and exhausts wealth in finite time. These problems can be corrected even if the expected real return can go negative. We also show that compounding may neutralize the preservation-of-capital constraint in optimization models, and we discuss formulations in which preservation of capital retains its bite.

Keywords: University endowment; Spending; Capital preservation; Destroying capital; Moving average; Smooth spending; Smooth consumption; Permanent contribution

JEL codes: G11; G23; L31

1 Introduction

Donors who wish to contribute to universities have a number of options depending on when they want their giving to have an impact. For example, donors wanting to have an immediate impact can contribute through annual giving, donors who want to have an impact for an intermediate time frame can give funds for a building, and donors who want to have a permanent impact can contribute to endowment. Since contributions to endowment are supposed to have a permanent impact, a university has a responsibility to make sure that its spending rule and investment strategy, taken together, preserve capital.[1] In other words, preservation of capital is viewed as a constraint on universities' choice of policy, and may be viewed as a way of requiring university leadership to preserve resources for the future without constraining leadership too much in the short term. This paper studies preservation of capital with a focus on existing practice. It is known that spending less than the expected return on investments is the wrong criterion for preservation of capital, although this criterion is still commonly used by practitioners. Our first major result is that the common practice of smoothing spending is destabilizing and implies that a unit of endowment reaches zero value in finite time, or at least the policy will have to be abandoned to avoid ruin. We also show how to modify the smoothing rule so it can preserve capital. Another major result is that the requirement of preservation of capital can be circumvented if general strategies are allowed. This is thanks to the miracle of compounding: setting aside an economically trivial amount of money (e.g. two cents) will grow without bound over time and satisfy the preservation-of-capital constraint. This result highlights the fact that the definition of preservation of capital is suitable for the sort of simple strategies commonly used in practice, but not for general strategies considered in dynamic optimization problems.

We start by reviewing an important result of Coiner (1990), which is important for practice and seems not to be understood well by practitioners. Coiner points out that

[1] Sedlacek and Jarvis (2010) discuss the legal foundation of this requirement in the Uniform Prudent Management of Institutional Funds Act.

choosing spending less than expected return is not the correct criterion for preserving capital. The intuition behind choosing spending less than expected return comes from the law of large numbers and the assertion that this means the value of a unit increases over sufficiently long time intervals. However, the law of large numbers applies to sums, not products (like the product over time of ratio of the value of a unit at the end versus the beginning of a period). As Coiner points out, the correct criterion looks for a positive expected real log return net of spending, which implies that the unspent remaining value of an initial investment increases without bound over time. This result should seem natural in light of the Kelley (1956) criterion. Samuelson's (1971) critique of the log-optimal policy does not apply, because preservation of capital is a constraint, not the objective function. We prove a straightforward generalization of Coiner's result that uses the ergodic theorem for covariance-stationary processes to show that capital is preserved if the expected log return net of spending is positive *on average*. We use this generalization to prove other results in the paper, including the result that a simple modification of the smoothing rule does preserve capital for some parameter values.

Our first main result shows that the common practice of smoothing spending is destabilizing. Smoothing of spending is supposed to prevent the damage done by large fluctuations in spending. This is a reasonable idea: sudden decreases in spending are disruptive, and sudden increases may be used carelessly. This is why many endowments (more than 70% according to NACUBO and commonfund (2017)) use a linear rule to smooth spending. Unfortunately, if the target spending rate is positive, capital is *never* preserved by the usual rule that moves only partway to a target spending level. The proof assumes the target is positive (but possibly very small) and that portfolio returns are risky and i.i.d. Intuitively, random fluctuations imply that sooner or later we will have bad luck in the risky investment, making the spending rate very large. Consequently, capital is depleted relatively more quickly than spending is reduced, and as a result the endowment ends up in a "death spiral" plunging to zero wealth in finite time. Our general result is proven in a risky model that uses the same mathematics, with a change of space and a change of time, that is used for boundary classification in stochastic differential equations. We also provide a simple riskless example that shows why the endowment is in serious trouble once the spending rate is high enough. If the initial spending rate is high enough, the proportional reduction in wealth from spending more than income is larger than the proportional reduction in spending from the smoothing rule, causing wealth to hit zero in finite time. We do not take the endowment reaching zero in finite time in these models too literally, since it is likely that at some point the endowment managers will realize they are on a path to doom and change the rule. Brown, Dimmock, Kang, and Weisbenner (2014) present evidence that endowments spend less than their stated policies after large drops in value but stay with their stated policies after increases.[2] Of course, it would be much better to have a sustainable rule that does not get endowments into terrible trouble in the first place.

Since smoothing is probably a good idea and the traditional smoothing rule never preserves capital, we propose a possible solution, a simple improved smoothing rule that includes a new term that changes spending to compensate for expected net overspending. Because the adjustment is for expected (rather than actual) overspending, the new rule's spending is still smooth (differentiable). For the rule with adjustment for expected overspending, we have a characterization of the parameter values for which capital is preserved. This characterization has a cap on the spending target that is the cap on spending absent smoothing times a factor less than one that depends on the investment return variance divided by the rate of adjustment towards the target rate. If the rate of adjustment towards the target is high and the variance is small, the spending rate will spend most of the time near the target, and the adjustment for expected shortfall is neutralized by the large mean reversion, and the cap on the target is almost the same as the cap on the spending rate absent smoothing. However, when

[2]Brown, Dimmock, Kang, and Weisbenner (2014) seems to suggest that deviating from the stated policy, a deviation they refer to as endowment hoarding, is a bad thing. In our view, the bad thing is the destabilizing policy itself, and deviating from the policy is necessary for survival when the endowment is on a trajectory to ruin.

the rate of adjustment is low and the variance is high, the spending rate will spend a lot of time away from the target spending rate. In this case, the target spending level must be smaller and the adjustment by the expected shortfall will move around more (because the spending rate is away from the target more). These features provide slack that allows the adjustment mechanism to work.

This paper investigates the conditions under which conventional investment and spending policies preserve capital. The focus is on the constraints endowments need to meet, in order to preserve capital as promised to the donors. This contrasts to the usual optimal investment approach taken by academics, which maximizes a utility function subject to constraints (see for example, Dybvig (1999), Gilbert and Hrdlicka (2015), or Campbell and Sigalov (forthcoming)). In general, practitioners find optimization models less useful than academics would hope, since it is difficult to incorporate all the considerations that are important in practice. Nonetheless, optimization models are useful benchmarks for thinking about new strategies. Although we do not solve any optimization models in this paper, we look at some implications of incorporating preservation of capital in these models. Our results suggest that the plausible definition of preservation of capital we are using, which is fine for the sorts of policies traditionally considered, can be subverted in an optimization problem that allows general strategies. An optimization model with general strategies will find an optimal manipulation to sidestep the rule, implying the constraint will either be irrelevant in an optimization problem or there will be no solution. Specifically, we prove that any utility level that can be obtained without the constraint on preservation of capital can also be approached arbitrarily closely with the constraint. Intuitively, this is because the constraint only imposes a condition on wealth in the limit as time increases, for which compounding can obtain a large value from a trivial investment. For example, the current college president may choose to spend all but two cents worth of the endowment before retiring, with a plan of modest spending afterwards. Theoretically, investing two cents in the riskless asset will grow without limit over time to satisfy the constraint on preservation of capital, without having any material effect on the current president's plans. Our proposed solution is to disallow investment and spending strategies that depend explicitly on the calendar date, or otherwise give "unfair" preference to current spending at some dates.

Preservation of capital is related to a recent literature on sustainable consumption (spending in our setting). Dybvig (1995) assumes strict sustainability: consumption cannot fall over time (or whose rate of decline is constrained), and Dybvig (1999) applies the model to managing an endowment. The idea is to apply the strict sustainability condition to spending that would be extremely damaging to interrupt, for example, salary for tenured faculty or funding for collecting a data series that would be much less valuable with gaps. Campbell and Sigalov (forthcoming) defines sustainability as either consumption or log consumption having a conditional expected change of zero.[3] Campbell and Martin (2021) defines sustainability as nondecreasing expected felicity of consumption over time, explicitly including examples in which there is no real riskless asset. Campbell and Martin (2021) focuses explicitly on society (possibly with population growth) rather than a single-agent choice problem like managing an endowment, and they have a survey of previous literature on sustainability from growth theory.

Whether sustainability and preservation of capital are nested depends on the model and how it is implemented. In practice, preservation of capital is likely (and traditionally) imposed on a unit of endowment, because it reflects a promise to donors who contribute endowment but not other doners. On the other hand, sustainability is likely imposed on some or all expenditures whatever the funding. This distinction is important in practice, because preservation of capital cannot be satisfied by the the anticipation of future contributions, but the distinction is not usually made in optimization models with constraints imposing preservation of capital or sustainability, which typically abstract from gifts after the outset. In the models with all funding

[3]Their constraint on consumption (or spending) c_t is in effect the same as $(\forall s < t)E_s[c_t] = c_s$ (or $(\forall s < t)E_s[\log(c_t)] = \log(c_s)$). Curiously, the analysis would be completely different than if they imposed the similar-looking constraint $(\forall t)E_0[c_t] = c_0$ (resp. $(\forall t)E_0[\log(c_t)] = \log(c_0)$).

at the outset, Dybvig (1999) implies preservation of capital if there is any protected spending, Campbell and Sigalov (forthcoming) do not preserve capital with either the linear or log constraint (but would if the expected log change were constrained to be positive instead of zero), and Campbell and Martin (2021) implies preservation of capital because of a model restriction to agents more risk averse than log, but otherwise would not. None of these models preserves capital if the current leadership can argue that there is no limit on current spending because future contributions will meet the sustainability conditions. There is another concern with sustainability in the setting of endowments. Applying sustainability to all spending would be too constraining in the short term (for which the leadership's incentives are not a big problem), ruling out lumpy spending for big projects (such as adding a wing to a hospital) and useful smoothing over time. In general, imposing sustainability on endowment spending is interesting (especially for categories of spending that would be painful to reduce), but not as a substitute for preservation of capital.

A lot of other research on endowments is complementary to our analysis: see Cejnek, Franz, Randl, and Stoughton (2013) for a survey and synthesis of many strands of the literature. Empirical work about the strategies and returns to endowments seem especially useful for understanding what sort of assumptions we should make about the return process. For example, Lerner, Schoar, and Wang (2008) have evidence that tilting into alternative assets has created more skewed portfolio returns, while Goetzmann and Oster (2015) provide evidence that the move into alternative assets may have been motivated partly by chasing past performance and partly by competition in the universities' market for students.

The rest of the paper is arranged as follows. Section 2 discusses background and documents some commonly used but problematic rules for preserving capital. As a preparation for proofs in the following sections, it also generalizes the result of Coiner (1990) to a case with ergodic covariance-stationary return mean, variance, and spending. Section 3 shows that traditional smoothing implies that capital is not preserved. We provide an improved smooth spending rule that preserves capital. Section 4 comes up with the condition for preserving capital with temporarily negative real risk-free rate. Section 5 discusses optimization models of spending and investment. Section 6 closes the paper.

2 Background

In the following subsection, we present a reasonable definition of preservation of capital that will be used in the paper.

2.1 Definition of Preservation of Capital

To characterize preservation of capital, we require a formal definition of what this means. Fortunately, most of our results will be robust to a range of reasonable choices for how we define preservation of capital. We stay close to the standard choice by practitioners, partly because the function of this paper is to evaluate their rules. The idea of a preservation-of-capital constraint is to protect the interests of future generations, given that current university leaders who might like to spend a lot now, without micromanaging short-term choices for which there is less of a conflict. For this reason, practitioners focus on the long-term evolution of real wealth from investments and net of spending out of an initial unit of endowment without putting constraints on spending in the short term.

We follow practitioners' reasonable standard that we should consider a unit of endowment, with a proportional change equaling the investment return less the spending rate, but not including any new contributions. Looking at a unit without credits is important because we are requiring a contribution to have a permanent impact. It is annual giving, not a permanent contribution to endowment, if we spend the entire contribution this year and replace it using future contributions. Including future contributions would be important for writing down an optimization problem, and in particular spending from future contributions should be included in an optimization problem's budget constraint, but not in constraints on preserving capital.

We let W_t be the real (inflation-adjusted) value of wealth in the unit at time t with spending S_t. Let r_t be the real rate of return and s_t be the spending rate at time t. We will not concern ourselves with valuation issues such as what price index to use or how to value illiquid assets, so that given the investment and spending policy for the endowment, the processes W_t, r_t, and s_t are well-defined. We also abstract from parameter uncertainty about the distribution of returns.

We will use the following definitions:

Definition 1 *Endowment wealth is said to be* preserved *if the real value of a unit W_t becomes arbitrarily large over time:* $\mathrm{plim}_{t\to\infty} W_t = \infty$.[4]

Definition 2 *Endowment wealth is said to be* destroyed *if the real value of a unit W_t vanishes over time:* $\mathrm{plim}_{t\to\infty} W_t = 0$.

We think of the definition of destroying capital as relatively conservative, since no reasonable rule for preserving capital would say that we are preserving capital if wealth is almost always close to 0 when t is large. Any other definition that implies we destroy capital in this case would also be sufficient for our results.

2.2 Problems in Common Institutional Policies

A traditional criterion for preservation of capital imposes a spending rate of no more than the average return on the endowment. Coiner (1990) and Ho, Mozes, and Greenfield (2020) have pointed out a flaw in this traditional criterion. Although this critique and the correction using logs is not new to this paper, it is worth repeating both because the correction does not seem to have been incorporated in practice and because we will use a generalized version of Coiner's result to prove some of our results. The traditional criterion is based on the intuition of using the law of large numbers, since

[4] As is conventional, plim indicates convergence in probability. By definition, $\mathrm{plim}_{t\uparrow\infty} W_t = \infty$ if for all $X > 0$, $\mathrm{prob}(W_t > X) \to 1$ as $t \to \infty$.

the average portfolio return over time will converge to the population mean. However, this is not the correct criterion. Terminal wealth is the initial wealth times the product over time of gross return less the spending rate:

$$W_t = W_0 \prod_{v=1}^{t} (1 + r_v - s_v), \tag{1}$$

where W_t is the wealth process of a unit, r_v is the return at time v, and s_v is the spending rate at time v. However, if the wealth relative $1 + r_v - s_v$ is i.i.d. and has mean bigger than 1, this does not imply that wealth must grow without bound over time, since the law of large numbers applies to sums, not products. If instead we take logs,

$$\log(W_t) = \log(W_0) \prod_{v=1}^{t} \log(1 + r_v - s_v), \tag{2}$$

and if $1 + r_v - s_v$ is i.i.d. with $E[\log(1 + r_v - s_v)] > 0$ and $\mathrm{var}(\log(1 + r_v - s_v)) < \infty$, the law of large numbers implies that $\log(W_t)$ increases without bound over time, and therefore so does W_t. Since log is a concave function, a standard application of Jensen's inequality proves that the corrected criterion is more stringent than the original criterion.[5]

As noted by Ho, Mozes, and Greenfield (2020), a standard calculation using Itô's lemma[6] can be used to derive the appropriate criterion in the usual continuous-time model with i.i.d. returns and lognormal stock prices. In particular, the local expected change in the log value of a unit must be positive, $\mu - \sigma^2/2 - s > 0$, where μ is the mean return, σ^2 its local variance, and s is the spending rate. If the portfolio's standard deviation of returns is 15%, the allowable spending rate is smaller by $\sigma^2/2 = 0.15^2/2 \approx 1\%$, which is a big difference given that spending rates are usually on the order of 3%

[5] Jensen's inequality says that $E[\log(1 + r_v - s_v)] < \log(E[1 + r_v - s_v])$. Therefore, $E[\log(1 + r_v - s_v)] > 0$ implies $\log(E[1 + r_v - s_v]) > 0$, which is equivalent to the traditional criterion $E[r_v - s_v] > 0$ since $\log(\cdot)$ is an increasing function and $\log(1) = 0$.

[6] If $dW_t = \mu W_t dt + \sigma W_t dZ_t - sW_t dt$, Itô's lemma says $E[d\log(W_t)] = (1/W_t)E[dW_t] + (-1/W_t^2)\sigma^2 W_t^2 dt/2 = \mu - s - \sigma^2/2$, where we have used the facts that $d\log(x)/dx = 1/x$ and $d^2\log(x)/dx^2 = -1/x^2$.

to 5%.

Although the deficiency of the traditional criterion is already known to academics, this criterion is still widely used in practice. For example, the spending policy statement of UCSD Foundation (2014) states that its objective is to "achieve an average total annual net return equivalent to the endowment spending rate adjusted for inflation." Moreover, the endowment of Henderson State University (2014) even employs a concrete example to illustrate its objective of achieving an inflation-adjusted average return equal to the spending rate: "Total return objective 7.00%, spending rate 4.00%, administration fee 1.50%, and inflation rate 1.50%." This criterion is also mentioned by Rice, Dimeo, and Porter (2012), which gives as a hypothetical example: "the primary objective of the Great State University Endowment fund is to preserve the purchasing power of the endowment after spending. This means that the Great State University Endowment must achieve, on average, an annual total rate of return equal to inflation plus actual spending."

A moving average smooth spending rule is another commonly used but one of our main results is that it is flawed rule. Specifically, instead of making spending strictly proportional to the size of the endowment, it is common to smooth spending using a moving-average (partial adjustment) rule to move from current spending towards a spending target. There is some economic sense to smoothing spending, since a sudden decrease in a budget can cause distress, while a sudden increase can invite waste. As a result, many endowments smooth their spending. For instance, several universities in the UC system use a smooth spending policy (Mercer Investment Consulting (2015)): UC Berkeley, UC Irvine, and UC Santa Cruz plan to spend about 4.5% of a twelve-quarter moving average market value of the endowment pool. Another example: Grinnell College Endowment (2014) states that endowment distribution is calculated as 4.0% of the twelve-quarter moving average endowment market value determined annually as of the December 31 immediately prior to the beginning of the fiscal year. According to the NACUBO and commonfund (2017) study of endowments, a "clear majority," 73% in fiscal year 2017, of respondents in their study "reported that they compute their spending by applying their policy spending rate to a moving average of endowment value." See also Acharya and Dimson (2007), Chapter 4, page 112. However, as we will see in the following section, the moving average rule is proved to destabilize the endowment, and we propose improved smoothing spending rules that preserves capital.

It is also implicit in our analysis that there is a degree of integrity in the endowment accounting process. For example, it would be improper for the university to borrow from the endowment and count the loan as an asset. This convention would misrepresent the value of the endowment and could be used to circumvent entirely any requirements for preservation of capital. Just spend whatever you want out of endowment, record the spending as a ten-year bullet loan, and when the loan matures roll it over into a new ten-year bullet loan. Using this device, we could spend the entire endowment without recording any spending at all. In our view, a university borrowing from its own endowment seems fraudulent since it misrepresents the value of the endowment, but we do not know how the law would view this.

2.3 Use Time-Average Log Returns and Spending Rates

As mentioned above, the criterion for preservation of capital in the i.i.d. case is $s < \mu - \sigma^2/2$. However, to investigate if more general and realistic spending rules, such as the moving average spending rule and the spending rules in an environment with stochastic interest rate, preserve capital, we need a generalized version of the log return criterion. As a preparation for the following sections, we generalize the log return criterion to a case with more general spending process and return process in this subsection. Let μ_t be the local mean return on the portfolio and let σ_t^2 be the local variance, and we will let μ_t and σ_t^2 follow some ergodic covariance-stationary processes. To get started, we review the definition of a covariance-stationary process in continuous time and a corresponding basic ergodic theorem.

Definition 3 *A stochastic process Y_t for $t \in (-\infty, \infty)$ is said to be covariance stationary (also called weakly stationary) if its mean is constant over time and autocovariances depend only on time difference. In notation, Y_t is stationary if there exists a mean M and autocovariance function $\Gamma(\Delta)$ such that, for all t and Δ, $\mathrm{E}[Y_t] = M$ and $\mathrm{cov}(Y_t, Y_{t+\Delta}) = \Gamma(\Delta)$.*

Theorem 1 *(Covariance Stationary Ergodic Theorem) A covariance stationary process Y_t with covariance function $\Gamma(\Delta)$ and mean M has a sample mean*

$$\frac{1}{T}\int_{t=0}^{T} Y_t dt$$

that tends to M in L^2 as T increases if and only if

$$\lim_{T\uparrow\infty} \frac{1}{T}\int_{\Delta=0}^{T} \Gamma(\Delta) d\Delta = 0.$$

Such a process is called an ergodic covariance-stationary process.

Proof: Immediate, from applying Shalizi and Kontorovich (2010), Theorem 298, to $X(t) \equiv Y_t - M$. ■

Given the ergodic covariance-stationary processes μ_t and σ_t^2, we can write a fairly general wealth dynamic as

$$dW_t = W_t\left(\mu_t dt + \sigma_t dZ\right) - S_t dt = \left(W_t \mu_t - S_t\right) dt + W_t \sigma_t dZ, \tag{3}$$

which implies that

$$W_t = W_0 \exp\left[\int_{v=0}^{t}\left(\mu_v - \frac{1}{2}\sigma_v^2 - s_v\right) dv + \int_{v=0}^{t} \sigma_v dZ_v\right]. \tag{4}$$

Then we can obtain a straightforward extension of the log criterion with the i.i.d. returns. Here is the formal result that gives conditions under which it suffices for capital to be preserved.

Theorem 2 *Let the wealth process be generated by (3), where* μ_t*,* σ_t^2*, and* s_t *satisfy* $\sigma_t > 0$ *and* $s_t > 0$*. If* μ_t*,* σ_t^2*, and* s_t *are all covariance stationary and ergodic, then capital is preserved if* $\mathrm{E}[s_t] < \mathrm{E}[\mu_t - \sigma_t^2/2]$ *and capital is destroyed if* $\mathrm{E}[s_t] > \mathrm{E}[\mu_t - \sigma_t^2/2]$ *(where the expectations do not depend on* t *because of stationarity).*

Proof. See Appendix A.1.

We have used a simple ergodic theorem to derive this result, but there are many obvious extensions, using one of the many other sufficient conditions for ergodicity. With a little structure on the stochastic processes for returns and spending, it is sufficient to preserve capital for spending to be less than the expected log return *on average*. This general result allows for time-varying parameters and also implies that we can keep spending in bad times with low interest rates and risk premium providing we also do not spend too much in good times.

Note that expected returns are not "sustainable" payout rates is related to that log-normally distributed returns are highly skewed with the mean being above the median. And it is also related to the "Kelly Rule": though expected returns converge with horizon, the wealth variance increases without bound, i.e., the fallacy of time diversification in an i.i.d world. This is why the maximal growth investment portfolio need not be optimal unless one has log utility (Samuelson 1971). Fortunately, this is not an issue for our paper, since the logarithm is in the preservation-of-capital constraint, not the objective function.

3 Preserving Capital with Smooth Spending

The commonly-used moving average rule tends to destabilize the endowment. We illustrate this with a riskless example for which an initial high spending rate sends the fund into a "death spiral" with the wealth going to zero for sure at a known finite time. Then we give a result for risky i.i.d. returns. When risky investment returns are bad, wealth goes down, but spending is slow to adjust so the spending rate goes up. At some point the fall in wealth becomes unstable because the adjustment is not fast enough to keep the spending rate from getting large as wealth (in the denominator) falls. In a risky investment environment, this scenario will play out sooner or later, and capital is always destroyed.

3.1 Traditional Moving Average Spending Rule: Riskless Case

We model the moving average rule used by practitioners as a partial adjustment model. The partial adjustment model is equivalent to a moving average model with exponentially declining weights. We assume smoothing in real terms, but this is not important, and indeed for reasonable cases smoothing in nominal terms can be rewritten as smoothing in real terms.[7]

$$dS_t = \kappa \left(\tau W_t - S_t\right) dt, \tag{5}$$

where τ is the target spending rate, and κ captures the adjustment speed. If the endowment only invests in a riskless bond with constant risk-free rate r, then the wealth process is given as

$$dW_t = rW_t dt - S_t dt. \tag{6}$$

We assume that if W_t reaches zero, then the endowment is shut down and W_t and S_t are both zero forever afterwards if wealth reaches zero. We will also assume $\tau < r$, which implies that spending at the target rate would preserve capital, so our policy has a fighting chance. We have the following result.

Theorem 3 *When the endowment only invests in a riskless asset, the moving average spending rule (5) does not preserve capital when the initial spending rate S_0/W_0 is sufficiently high. Specifically, given the dynamic (5) and (6), wealth W_t reaches* 0*, if*

[7] For example, if inflation is always at a rate ι and the moving average rule is applied in nominal terms, (5) should be replaced by $dS_t = \kappa\left(\tau W_t - S_t\right) dt - \iota S_t dt$, where the new final term gives the rate of reduction of real spending due to inflation. However, this expression is identical to (5) if we redefine κ to be $\kappa + \iota$ and we redefine τ to be $\tau\kappa/(\kappa + \iota)$. This redefinition preserves $\kappa > 0$ and $\tau > 0$ provided $\iota > 0$ (or at least $\iota > -\kappa$), so the same analysis applies.

S_0/W_0 is large enough, in finite time t^*, and

$$t^* = \frac{1}{\lambda_1 - \lambda_2} \log \left(-\frac{K_2}{K_1} \right),$$

where

$$K_1 = \frac{W_0 \left(r - \lambda_2 \right) - S_0}{\lambda_1 - \lambda_2} \text{ and } K_2 = \frac{W_0 \left(\lambda_1 - r \right) + S_0}{\lambda_1 - \lambda_2},$$

and $\lambda_2 < 0 < \lambda_1$ is given by

$$\lambda_1 = \frac{r - \kappa + \sqrt{\left(\kappa - r \right)^2 - 4\kappa \left(\tau - r \right)}}{2} \text{ and } \lambda_2 = \frac{r - \kappa - \sqrt{\left(\kappa + r \right)^2 - 4\kappa\tau}}{2}.$$

Proof. See Appendix A.2.

If the endowment starts with high spending under the moving average rule, capital will be wiped out quickly. Given a high initial spending rate, the value of a unit declines proportionately more (due to the shortfall of interest covering spending) than spending (due to the moving average rule). As the ratio of wealth to spending falls, this effect accelerates and wealth converges to zero in a "death spiral."

3.2 Traditional Moving Average Spending Rule: Risky Case

We have just seen that if the initial spending rate is high enough, an endowment making a riskless investment and smoothing towards any positive target spending rate will destroy capital. In this subsection, we show that an endowment smoothing towards a target spending rate and risky portfolio strategy will destroy capital for any initial spending rate. The intuition is that the random portfolio returns will lead us sooner or later into a situation with high spending that will deplete the assets.

To model wealth, we have to make an assumption about the portfolio returns. The portfolio choices of endowments in practice are not usually linked dynamically to the current spending rate.[8] Usually, the percentage allocations to different asset classes have fixed target values or ranges. As a result, it is a reasonable approximation to take the endowment return process as exogenous when we model the spending rule. Therefore, we let

$$\begin{aligned} dW_t &= W_t\left(\mu_t dt + \sigma_t dZ\right) - S_t dt \\ &= \left(W_t\mu_t - S_t\right) dt + W_t\sigma_t dZ, \end{aligned} \tag{7}$$

where μ_t and σ_t are exogenous processes, so long as wealth stays positive. Also assume that zero is an absorbing barrier for wealth, that is, if W_t reaches zero, then the unit is "shut down" and W_t and S_t are both zero forever afterwards if wealth ever reaches zero. We have the following result.

Theorem 4 *Suppose an endowment uses the moving average spending rule (5) with positive target spending rate τ, no matter how small, and the i.i.d. investment process (7) with fixed $\mu_t = \mu > 0$ and $\sigma_t = \sigma > 0$. Then the value of a unit hits zero in finite time (almost surely) and therefore capital is always destroyed according to Definition 2.*

Sketch of proof: Given the joint dynamics of wealth and spending, we can write the dynamics of wealth over spending (which is Markovian). Then find a function F of the variable W_t/S_t such that $F(W_t/S_t)$ is a local martingale (by deriving the dynamics of $F(W_t/S_t)$ using Itô's lemma, and set the drift term equal to zero). We prove that $F(0)$ is finite, $F(\infty) = \infty$ and F is one-to-one on $[0, \infty]$. Since $F(W_t/S_t)$ is a continuous local martingale, we can change time to a Wiener process with constant variance per unit time. We use the known properties of the first-hitting problem with constant variance and the properties of the time change (using the local variance of $F(W_t/S_t)$) to show that $F(W_t/S_t)$ hits $F(0)$ in finite time, at which time therefore W_t/S_t hits zero. See Appendix A.3 for the detailed proof. ■

[8] Ideally, investment mix should depend on the spending rate, see Dybvig (1999), but the assumption of independence is closer to practice and is good for illustrating our concerns.

Recall that in the riskless case, a unit enters a death spiral to zero if initial spending is high enough, since the proportional decrease in spending does not keep up with the proportional decrease in wealth. In the stochastic case, sufficiently bad luck in investments over a short time depletes the wealth, increasing the spending rate to a high level, starting a death spiral. Subsequent good luck can save the endowment, but sooner or later the endowment will have sufficiently bad luck starting a death spiral the endowment does not recover from.

We needn't take wealth hitting zero too literally. In practice, it is probable that the strategy would be abandoned once it is clear that the endowment is in deep trouble. The point is that we should not be following a strategy that will require this sort of desperate rescue sooner or later, or to put it another way, we are not really following this policy.

3.3 A Smooth Spending Rule that Can Preserve Capital

We have seen that traditional moving average rules are destabilizing. Since smoothing is desirable, we now propose an improved version of the moving average rule that does preserve capital for some parameter values. The problem with the moving average rule is that when spending is high, the rate of reduction in spending from the moving average rule is less than the rate of reduction in wealth from spending more than the geometric average return on assets. Specifically, for any smooth (time differentiable) spending process S_t, Itô's lemma implies that

$$d\log\left(\frac{S_t}{W_t}\right) = S_t' dt - \left(\mu_t - \frac{\sigma_t^2}{2} - \frac{S_t}{W_t}\right) dt - \sigma_t dZ_t.$$

When the spending rate S_t/W_t is high, any decrease in spending S_t' has to overcome the term $S_t/W_t - (\mu_t - \sigma_t^2/2)$. The reason that traditional smoothing destroys capital when μ_t and σ_t are constant is that the term $S_t/W_t - (\mu - \sigma^2/2)$ dominates when the spending rate is large so that the spending rate gets larger and larger until all wealth is exhausted. This motivates changing the original smoothing rule (5) to the improved smoothing rule

$$dS_t = S_t \left(\underbrace{\kappa \left(\log \tau - \log \left(\frac{S_t}{W_t} \right) \right)}_{\text{Smooth spending with target } \tau} + \underbrace{\mu_t - \sigma_t^2/2 - \frac{S_t}{W_t}}_{\text{Adjustment for } mean \text{ over-spending}} \right) dt, \tag{8}$$

where we continue to assume the wealth process (7). This improved spending rule is smooth (S_t is a differentiable function of time), unlike a fixed spending rate (for which $S_t = \tau W_t$ is not differentiable because W_t is not). Unlike the original smoothing rule (5), this rule does not always destroy capital. As for the spending rules in Section 2, whether this improved rule preserves capital depends on the parameters.

With the proposed spending rule (8), we can prove the following theorem.

Theorem 5 *Suppose the endowment wealth process follows (7) where $\sigma_t = \sigma > 0$ is constant and μ_t is covariance stationary and ergodic. Then the smooth spending rule given by (8) preserves capital if the following condition holds:*

$$\tau < \exp\left(-\frac{\sigma^2}{4\kappa} \right) \left[\mathrm{E}[\mu_t] - \frac{\sigma^2}{2} \right], \tag{9}$$

while capital is destroyed if the inequality is reversed.

Sketch of proof: Given the spending and wealth dynamics, $\log(s_t) = \log(S_t/W_t)$ is a stationary Gaussian process and it can be derived that s_t is an ergodic covariance-stationary process (an Ornstein-Uhlenbeck velocity process). Then the result follows from Theorem 2. See Appendix A.4 for the detailed proof. ■

The condition (9) means that the expected log growth rate of investment has to be larger than the long-term average spending rate $\mathrm{E}[S_t/W_t] = \tau \exp\left(\sigma^2/\left(4\kappa\right)\right)$. We can compute the long-term average because we know the exact process followed by S_t/W_t. This is why we need σ constant. We don't need μ_t to be constant because μ_t cancels when we compute the process for s_t. If the speed κ of mean-reversion is very large and

σ is small, then the spending rate is usually very close to the target spending rate τ, which is why (9) converges as κ increases or σ decreases to the formula $\tau < \mathrm{E}[\mu_t] - \sigma^2/2$ for a fixed spending rate τ, but otherwise the rule is more stringent.

4 Real and Nominal Rates

So far, we have been assuming that the interest rates are expressed in real terms. In this section we consider an example in which we explicitly separate real and nominal interest rates. This naturally includes a consideration of what to do when real rates are negative, currently a topic of significant interest to practitioners.

These calculations by practitioners are done in real terms (as they should be). An interest rate environment like the current one where inflation exceeds the nominal rate is a special challenge. The endowment never preserves capital if the expected real rate of return on the portfolio is always negative. If the real riskfree rate is negative, it is possible but not necessarily true, that investing in an asset with a positive risk premium will generate the positive real returns we need to preserve capital in the long-run. However, under some conditions, capital can be preserved even if the real expected rate of return on the portfolio is temporarily negative. This subsection models a temporarily negative real rate and provides the conditions needed for preserving capital by employing the results of Theorem 2.

Let the nominal interest rate η_t and the inflation rate ι_t be modeled by some Itô processes.[9] Hence, the stock price follows a diffusion process as

$$\frac{dP_t}{P_t} = (\eta_t - \iota_t + \pi_t)\, dt + \sigma_t dZ_t, \tag{10}$$

where π_t is the risk premium. With a fixed proportion θ in the stock, the wealth process

[9] As an alternative modeling strategy, we might think that the inflation rate is not known (as in Section 7 of Cox, Ingersoll, and Ross (1985)), and that there may or may not be a real riskless asset that is different from the nominal riskless asset. Conceptually, the analysis for these cases should be similar to what is given here.

follows,

$$
\begin{aligned}
dW_t &= (\eta_t - \iota_t) W_t dt + W_t \theta (\sigma_t dZ_t + \pi_t dt) - S_t dt \\
&= W_t ((\eta_t - \iota_t + \theta \pi_t) dt + \theta \sigma_t dZ_t) - S_t dt.
\end{aligned}
$$

Theorem 2 implies the following theorem:

Theorem 6 *Assume the reinvested stock price follows (10), the endowment has a constant proportion θ in stock, and the spending rate s_t, nominal rate η_t, inflation rate ι_t, risk premium π_t, and variance σ_t^2 are ergodic covariance-stationary processes. Then the endowment preserves capital if*

$$
\mathrm{E}\left[\eta_t - \iota_t + \theta \pi_t - \frac{\theta^2 \sigma_t^2}{2}\right] > \mathrm{E}\left[s_t\right], \tag{11}
$$

and destroys capital if the inequality is reversed.

Now we can analyze preservation of capital in an example of spending rule with possibly negative real interest rates.

Example 1 (Preservation of capital with temporarily negative real rate) Let the nominal interest rate follows a Cox, Ingersoll, and Ross (1985) (CIR) square-root process, i.e.,

$$
d\eta_t = a(b - \eta_t) dt + c\sqrt{\eta_t} dZ_t^\eta, \tag{12}
$$

where $a > 0$ is a constant adjustment speed, $b > 0$ is the long-term mean of the nominal interest rate, and c is the local volatility of the nominal rate. Further, let the inflation rate be $\iota_t = \gamma \eta_t + h_t$ where $\gamma \in [0, 1]$ and h_t also follows a CIR square-root process

$$
dh_t = a_h (b_h - h_t) dt + c_h \sqrt{h_t} dZ_t^h, \tag{13}
$$

where $a_h > 0$ and $b_h > 0$. This structure has the reasonable feature that the interest rate and the inflation rate tend to move together but are not perfectly correlated;

maybe $\gamma = 0.8$ is a reasonable value. We take $c^2 \leq 2ab$ and $c_h^2 \leq 2a_h b_h$ so η_t and ι_t never reach zero. We also assume that Z, Z^η, and Z^h are drawn independently. Note that η_t and h_t (and therefore ι_t) are ergodic covariance-stationary processes.[10]

Consider a fixed proportional allocation θ to the stock with spending according to the improved smoothing rule (8), which now can be written as

$$dS_t = S_t \left(\kappa \left(\log \tau - \log \left(\frac{S_t}{W_t} \right) \right) + \eta_t - \iota_t + \theta \pi_t - \frac{\theta^2 \sigma_t^2}{2} - \frac{S_t}{W_t} \right) dt.$$

Take the reinvested stock process to be (10) where $\sigma_t = \sigma > 0$ and the risk premium π_t is an ergodic covariance-stationary process. Then $\mu_t = \eta_t - \iota_t + \theta\pi_t$ and the local variance of wealth is $(\theta\sigma)^2$, which, by the results in Theorem 5, implies capital is preserved if $\mathrm{E}[\eta_t - \iota_t + \theta\pi_t - \theta^2\sigma^2/2] > \tau \exp(\theta^2\sigma^2/(4\kappa))$. Now, $\mathrm{E}[\eta_t] = b$ and $\mathrm{E}[\iota] = \gamma b + b_h$, so capital is preserved if $(1-\gamma)b - b_h + \theta\mathrm{E}[\pi_t] - \theta^2\sigma^2/2 > \tau\exp(\theta^2\sigma^2/(4\kappa))$. The real rate in this example, $\eta_t - \iota_t = (1-\gamma)\eta_t - h_t$ can be negative since η_t and h_t evolve independently and η_t can be very close to zero and h_t can be arbitrarily large. The average real rate, $(1-\gamma)b - b_h$, can be positive or negative. Even if it is negative, the average log risk premium on equities, $\theta\mathrm{E}[\pi_t] - \theta^2\sigma^2/2$, may be large enough to preserve capital if τ and the other factor on the right-hand-side of the inequality are not too large.

More generally, there is nothing to guarantee that it is even feasible to preserve capital. If the current environment with small nominal interest rate and larger inflation is the "new normal," then the mean real rate of return from holding riskless assets could be negative, and must be overcome by the log risk premium from the risky position in our portfolio. Since spending cannot be negative, it is not feasible to preserve capital if the mean log return from riskless assets is negative and of a larger magnitude than the largest available log risk premium. Probably this is not an issue now, since the

[10] From the comment in CIR after (20), the steady-state variance of η_t is $\mathrm{var}(\eta_t) = c^2 b/2a < \infty$ and therefore from the first line of (19) $\mathrm{cov}(\eta_s, \eta_t) = \mathrm{var}(\eta_t)\exp(-a|s-t|)$ which implies η_t is covariance stationary. Ergodicity follows from Theorem 1, since the limit is zero. The same argument applies to h_t.

historical risk premium is pretty large and the historical mean return on riskless assets is positive, but this depends on our beliefs of what will happen in the future.

5 Optimization Models

We have been emphasizing preservation of capital as a constraint facing by the universities. The traditional practice by endowments postulates a candidate portfolio strategy and spending rule, followed by a check of what parameter values, e.g., spending rate target and portfolio weights, are consistent with preservation of capital. Alternatively, we can impose preservation of capital as a constraint in an optimization problem. Unfortunately, simply imposing preservation of capital as a constraint in a standard optimization problem does not work. It is reasonable to apply Definition 1 to traditional strategies commonly adopted by practitioners, but an optimization problem will find a strategy that "optimally" subverts Definition 1 and renders it toothless. Intuitively, it is possible to set aside an arbitrarily small amount of the endowment that will grow without bound over the time and do as we like with the rest. We investigate this using the following Problem 1.

Problem 1 *Given the initial wealth W_0, choose adapted portfolio process $\{\theta_t\}_{t=0}^{\infty}$, adapted spending process $\{S_t\}_{t=0}^{\infty}$ and wealth process $\{W_t\}_{t=0}^{\infty}$ to maximize expected utility,*

$$
\begin{aligned}
&\mathrm{E}\left[\int_{t=0}^{\infty} D_t u\left(S_t\right) dt\right] \\
s.t.\ \ & dW_t = rW_t dt + \theta_t\left(\left(\mu - r\right) dt + \sigma dZ_t\right) - S_t dt, \\
& (\forall t) W_t \geq 0,\ \ \textit{and} \\
& \operatorname*{plim}_{t\to\infty} W_t = \infty,
\end{aligned}
\tag{14}
$$

where the utility function $u : \Re_+ \to \Re$ is concave and increasing. It is assumed that

$\mu - r$, σ, and r are all positive and the utility discount factor $D_t \geq 0$, and

$$0 < \int_{t=0}^{\infty} D_t dt < \infty.$$

The constraint (14) is preservation of capital according to Definition 1. The functional form of the objective function is flexible enough to accommodate the short-term orientation of a college president who does not value spending beyond the end of his term. For example, if the president is confident of retiring by time T, then perhaps $D_t = 0$ for all $t > T$.

The weakness of the constraint is that it concerns only the infinite limit, but does not restrict what happens at intermediate dates. And, due to the miracle of compounding, it only takes a small amount of money set aside to satisfy the condition that a unit grows without limit over time. Intuitively, we can put almost 100% of the endowment in our favorite strategy absent the constraint on preservation of capital and two cents in a strategy that preserves capital, to achieve almost the same utility as our favorite strategy. In this way, we can make the impact of the constraint on both our strategy and our utility negligible. Here is the formal statement:

Theorem 7 *Let S_t^*, θ_t^*, and W_t^* be the feasible spending, risky asset portfolio and wealth processes with finite value for Problem 1 without the preservation-of-capital constraint (14). Then, if $r > 0$, the supremum in Problem 1 is at least the value of following this strategy. Specifically, there exists a sequence (θ_t^m, S_t^m) of feasible policies satisfying the preservation-of-capital constraint (14), such that*

$$\lim_{m \to \infty} \mathrm{E}\left[\int_{t=0}^{\infty} D_t u\left(S_t^m\right) dt\right] \geq \mathrm{E}\left[\int_{t=0}^{\infty} D_t u\left(S_t^*\right) dt\right].$$

Proof. See Appendix A.5.

Note: as should be clear from the proof, most of the special structure of Problem 1 is not needed. All we need is that there is some asset or investment strategy that can (through the miracle of compounding) convert a trivial amount of capital today into an unbounded sum over time. This is a very mild condition: if it is not satisfied, then Problem 1 is not feasible in the first place.

Theorem 7 implies that although the preserving capital according to Definition 1 does restrict traditional strategies, it does not restrict general strategies, and optimization will sidestep the constraint. In the following subsection, we discuss approaches to formulating an optimization problem in which preservation of capital is not subverted.

5.1 Preservation of Capital in Optimization Models

Theorem 7 highlights that the traditional type of definition of preservation of capital in Definition 1 is not adequate for all strategies. There are two natural solutions to this problem. One is to come up with a more stringent definition of preservation of capital. Unfortunately, a more stringent definition tends to make preservation of capital infeasible. For example, if there are time periods when the real rate of interest is sufficiently negative given the maximum Sharpe ratio, then having a nonnegative expected log rate of growth is not feasible given that spending cannot be negative. The other natural solution (our preferred choice) is to restrict strategies to rule out explicit or implicit preference for spending now over spending in the future. For example, we could require dynamics of spending to depend only on current wealth and spending but not the calendar date. This is our preferred solution to the problem highlighted in Theorem 7: unlike making the definition of preservation of capital stricter, restricting attention to a set of "fair" strategies will not tend to make the optimization problem infeasible.

In terms of the first approach, we can think of a lot of stronger constraints, but it is hard to find a definition that is strong enough to give a significant constraint on behavior but is still feasible. A very strong definition is that wealth can never fall in real terms. This will definitely preserve capital if it is followed, but it precludes taking advantage of the market risk premium. More troubling is that it requires a asset or portfolio whose realized real is never negative if we want to have nonnegative consumption. TIPs or other inflation-protected government bonds may be close enough to being riskless in real terms, but their real returns are sometimes negative. Consequently, given that spending is nonnegative, it is infeasible to preserve capital in this very strong sense. A less stringent definition is that the real expected change in log wealth can never be negative. This does not require there is a real riskless asset, and perhaps more importantly, it is consistent with a negative real rate so long as the risk premium on another portfolio (perhaps the market) is high enough and its volatility not too high. Even this constraint may be infeasible if the real interest can go very negative and the market log risk premium does not always go up enough at the same time. For example, there is no strategy in Example 3 that would preserve capital in this stronger sense, because the real interest rate can be arbitrarily negative without an offsetting increase in log return from tilting into equities.

To make the wealth constraint feasible and more effective in preservation of capital, we can impose the drawdown constraint introduced by Grossman and Zhou (1993), which requires that wealth can never fall below a certain percentage of the previous maximum of wealth, i.e., for some given $\beta \in (0,1)$ and for all times t,

$$W_t \geq \beta \sup_{s \leq t} W_s.$$

The drawdown constraint carries a strong sense of preservation of capital, since it adds requirements on intermediate wealth, but it still allows for temporary losses because of general market conditions and taking advantage of the market risk premium. However, even this constraint may not be feasible if the real interest rate can go negative and stay there for a long time. In particular, given $\beta \in (0,1)$, there is no feasible strategy in Example 3 satisfying the drawdown constraint: for any strategy and some distant future date, we might have a run of negative real rates and bad luck on any stock tilt that would cause us to violate the drawdown constraint at that date.

The second approach requires that the portfolio choice and wealth dynamics are not allowed to give special preference to current spending. An explicit preference for current spending might follow one rule from now until 20 years from now, and another rule afterwards. An implicit preference might be a rule specifying higher spending during a situation (say high unemployment) that exists at the time the rule is enacted. Although this rule does not explicitly mention the calendar date, it gives preference to spending at the time the rule is put in place.

As an example of an optimization problem that restricts to non-preferential strategies, consider the following problem featuring preference for smooth spending (as in Section 3):

Problem 2 *Given the initial wealth W_0 and initial spending S_0, choose a proportional portfolio rule $\vartheta : \Re \to \Re$, an adapted proportional spending drift rule $\delta : \Re \to \Re$, and adapted wealth process $\{W_t\}_{t=0}^{\infty}$, to maximize expected utility,*

$$\begin{aligned}
&\mathrm{E}\left[\int_{t=0}^{\infty} e^{-\rho t}\frac{{S_t}^{1-R}}{1-R}dt\right], \\
s.t.\ \ & dW_t = rW_t dt + \vartheta(S_t/W_t)W_t\left((\mu - r)\,dt + \sigma dZ_t\right) - S_t dt - k(\delta(S_t/W_t))^2 S_t dt, \\
& \frac{dS_t}{S_t} = \delta(S_t/W_t)dt, \\
& (\forall t) W_t \geq 0,\ \ \textit{and} \\
& \operatorname*{plim}_{t\uparrow\infty}\frac{1}{t}\log\left(\frac{W_t}{W_0}\right) \geq \epsilon,
\end{aligned} \tag{15}$$

where we have various constants: $\rho > 0$ is the pure rate of time preference, $R > 0$, $R \neq 1$ is the constant relative risk aversion, $k > 0$ is the aversion to changing spending, r is the real rate of interest, and $\epsilon > 0$ is some small number.

Compared to Problem 1, a major difference is that the dynamics for portfolio choice and spending can only depend on current wealth and spending and not explicitly on time or recent random realizations that might be used to select for the current time. This restriction is introduced formally by not allowing the proportional portfolio choice ϑ and the proportional change of spending δ to be general processes and in particular restricting them to depend on the current spending rate and not on the calendar date. This problem also incorporates preference for smooth spending via the parameter k. The constraint with the small number $\epsilon > 0$ is a variant of the preservation-of-capital constraint that avoids a possible closure problem.[11] Problem 2 is homogeneous so the solution for δ, ϑ, and W_t/W_0 do not depend on initial wealth W_0 but only on the initial spending rate S_0/W_0.

Solving Problem 2 is feasible, at least numerically, but is beyond the scope of this paper. It might be nice to add other realistic features including contributions to annual spending and additions to endowment which may not affect the constraint for preservation of capital but would affect the objective function. Adding more state variables (e.g. the real interest rate) makes it a challenge for defining what sort of policy is just responding naturally to economic conditions versus what sort of policy is just designed to justify more spending now in an unreasonable way. In general, this approach offers both challenges and opportunities for future research.

6 Conclusion

Commonly-used rules for managing endowments that are supposed to preserve capital actually do not preserve capital. In addition to Coiner (1990)'s observation about spending less than expected return not being strict enough, the moving-average smooth spending rule tends to cause a "death spiral" of endowment and to destabilize capital, when returns are random. We provide an improved rule that smooths spending but in a way that preserves capital for appropriate choice of parameter values. We also discuss how to preserve capital in an environment with negative real rate as a special challenge like the current one.

In this paper, the focus was from the practitioner's lens and not based on optimiza-

[11] For example, suppose we choose spending and portfolio choice in the lognormal model with constant μ, σ, and s discussed after (4). Capital is preserved if $s < \mu - \sigma^2/2$; an agent who wants to spend more can improve on any feasible solution by increasing s towards $\mu - \sigma^2/2$. In this setting, preservation of capital (Definition 1) is equivalent to the existence of $\epsilon > 0$ such that (15) is satisfied.

tion, which is less useful for practitioners than we would hope. Nonetheless, solving optimization models may be more useful than most practitioners think, since seeing the optimal policy in an ideal setting can inform our intuition about what to consider in the more complex reality. We showed that it is not so obvious how to include preservation of capital as a constraint in an optimization model, since a natural optimization model can bypass this requirement entirely. Ruling out strategies that give "unfair" preference to current conditions can be a way forward, but requires more development.

We hope our results will help universities to do a better job managing their endowments.

References

[1] Acharya, Shanta and Elroy Dimson (2007). Endowment asset management: investment strategies in Oxford and Cambridge. Publisher: Oxford University Press.

[2] Arnold, Ludwig (1974). Stochastic Differential Equations: Theory and Applications. New York: Wiley.

[3] Brown, Jeffrey R., Stephen G. Dimmock, Jun-Koo Kang, and Scott J. Weisbenner (2014). How University Endowments Respond to Financial Market Shocks: Evidence and Implications. American Economic Review 104(3), 931–962.

[4] Campbell, John Y., and Ian W. R. Martin (2021). Sustainability in a Risky Environment. Mimeo, Harvard and LSE.

[5] Campbell, John Y., and Roman Sigalov (forthcoming). Portfolio Choice with Sustainable Spending: A Model of Reaching for Yield. Journal of Financial Economics.

[6] Cejnek, Georg, Richard Franz, Otto Randl, and Neil Stoughton (2014). A Survey of University Endowment Management Research. Journal of Investment Management 12(3), 90–117.

[7] Coiner, Michael (1990). The Lognormality of University Endowment in the Far Future and its Implications. Journal of Education Review 9(2), 157–161.

[8] Cox, John C., Ingersoll, Jr., Jonathon E., and Ross, Stephen A. (1985). A Theory of the Term Structure of Interest Rates. Econometrica 53(2), 385–407.

[9] Dybvig, Philip H. (1995). Duesenberry's Ratcheting of Consumption: Optimal Dynamic Consumption and Investment Given Intolerance for Any Decline in Standard of Living. The Review of Economic Studies 62(2), 287–313.

[10] Dybvig, Philip H. (1999). Using Asset Allocation to Protect Spending. Financial Analysts Journal, January-February, 49–62.

[11] Gilbert, Thomas and Christopher Hrdlicka (2015). Why Are University Endowments Large and Risky? Review of Financial Studies 28(9), 2643–2686.

[12] Goetzmann, William and Sharon Oster (2015). Competition Among University Endowments, in How the Financial Crisis and Great Recession Affected Higher Education. Eds. J. Brown and C. Hoxby. Chicago: NBER and University of Chicago Press.

[13] Grinnell College Endowment (2014). Grinnell College Endowment Spending Policy. Approved by the Foundation Board of Trustees in February 2014.

[14] Grossman, Sanford, and Zhongquan Zhou (1993). Optimal Investment Strategies for Controlling Drawdowns. Mathematical Finance 3, 241–276.

[15] Halmos, P. R. (1956). Lectures on Ergodic Theory. Math. Soc. Japan.

[16] Henderson State University Foundation (2014). Henderson State University Foundation Investment Policy Statement. Approved by the Foundation Board of Directors on December 18th, 2013; Amended on March 12th, 2014.

[17] Ho, Gregory P., Haim A. Mozes, and Pavel Greenfield (2010), The Sustainability of Endowment Spending Levels: A Wake-up Call for University Endowments, Journal of Portfolio Management 37(1), 133–146.

[18] Lerner, Josh, Antoinette Schoar, and Jialan Wang (2008), Secrets of the Academy: The Drivers of University Endowment Success, Journal of Economic Perspectives 22(3), 207–222.

[19] Mercer Investment Consulting, Inc. (2015). UC Annual Endowment Report Fiscal Year 2014-2015. Mercer Investment Consulting, Inc.

[20] NACUBO and commonfund (2017). NACUBO-commonfund Study of Endowments (NCSE), 2017.

[21] von Neumann, J. (1932). Proof of the Quasi-ergodic Hypothesis. Proceedings of the National Academy of Sciences, USA 18, 70-82.

[22] Rice, M., R. A. DiMeo, and M. Porter (2012). Nonprofit Asset Management: Effective Investment Strategies and Oversight. Publisher: John Wiley Sons.

[23] Rogers, L. C. G., and David Williams (2000), Itô Calculus, vol. 2 of Diffusions, Markov Processes, and Martingales. Cambridge, England: Cambridge University Press, 2nd edn.

[24] Samuelson, P. A. (1971). The "Fallacy" of Maximizing the Geometric Mean in Long Sequences of Investing or Gambling, Proceedings of the National Academy of Sciences of the USA 68, 2493–2496.

[25] Sedlacek, Verne O. and William F. Jarvis (2010). Endowment Spending: Building a Stronger Policy Framework. White paper, commonfund Institute.

[26] Shalizi, Cosma Rohilla and Aryeh Kontorovich (2010). Almost None of the Theory of Stochastic Processes. Manuscript, Carnegie Mellon, downloaded November 26, 2018 from http://www.stat.cmu.edu/%7Ecshalizi/almost-none/v0.1.1/almost-none.pdf.

[27] UC San Diego Foundation (2014). UC San Diego Foundation Endowment Investment And Spending Policy. Approved by the Board of Trustees, Effective December 12th, 2014.

A Appendix

A.1 Proof of Theorem 2

First we provide a lemma that will be used in the proof. It shows that σ_t^2 ergodic covariance-stationary implies that the long-term average of the changes $\sigma_t dZ_t$ is zero.

Lemma 1 *Suppose that σ_t^2 is an ergodic covariance-stationary process. Then*

$$\frac{1}{T}\int_{t=0}^{T}\sigma_t dZ_t$$

converges to 0 *in* L^2 *as* $T \to \infty$.

Proof: We first derive a bound on the variance of the integral of σ_t^2 from its ergodic property. By Theorem 1, $(1/T)\int_{t=0}^{T}\sigma_t^2 dt$ converges in L^2 to $\mathrm{E}[\sigma^2]$, which implies that for all $\varepsilon > 0$, there exists T^* such that for all $T > T^*$, $(1/T)\int_{t=0}^{T}\sigma_t^2 dt < \mathrm{E}[\sigma_t^2] + \varepsilon$, or equivalently, $\int_{t=0}^{T}\sigma_t^2 dt < (\mathrm{E}[\sigma_t^2] + \varepsilon)T$. Fix any $\varepsilon > 0$. Then for all T larger than the corresponding T^*,

$$\begin{aligned}
\mathrm{E}\left[\left(\frac{1}{T}\int_{t=0}^{T}\sigma_t dZ_t\right)^2\right] &= \mathrm{var}\left(\frac{1}{T}\int_{t=0}^{T}\sigma_t dZ_t\right) + \left(\mathrm{E}\left[\frac{1}{T}\int_{t=0}^{T}\sigma_t dZ_t\right]\right)^2 \\
&= \frac{1}{T^2}\int_{t=0}^{T}\sigma_t^2 dt \\
&< \frac{1}{T^2}(\mathrm{E}[\sigma_t^2] + \varepsilon)T \\
&\to 0 \quad \text{as } T \to \infty,
\end{aligned}$$

where the second and third steps follow because a driftless Itô integral with an L^2 integrand has mean zero and variance the squared L^2 norm of the integrand (see Arnold (1974, Theorem 4.4.14(e))). ■

Proof of Theorem 2: From (4), we have

$$\log\left(\frac{W_T}{W_0}\right) = \int_{t=0}^{T}\left(\mu_t - \frac{1}{2}\sigma_t^2 - s_t\right)dt + \int_{t=0}^{T}\sigma_t dZ_t$$

Applying Theorem 1 to each of μ_t, σ_t^2, and s_t and applying Lemma 1, we conclude that

$$
\begin{aligned}
\frac{1}{T}\log\left(\frac{W_T}{W_0}\right) &= \frac{1}{T}\int_{t=0}^{T}\mu_t dt - \frac{1}{T}\int_{t=0}^{T}\frac{\sigma_t^2}{2}dt - \frac{1}{T}\int_{t=0}^{T}s_t dt + \frac{1}{T}\int_{t=0}^{T}\sigma_t dZ_t \\
&\to \mathrm{E}[\mu_t] - \mathrm{E}[\sigma_t^2/2] - \mathrm{E}[s_t] \quad \text{as } T\to\infty
\end{aligned}
$$

in L^2 and therefore in probability. Consequently, $\operatorname{plim}_{T\to\infty} W_T = +\infty$ (preserving capital) if $\mathrm{E}[\mu_t] - \mathrm{E}[\sigma_t^2/2] - \mathrm{E}[s_t] > 0$ and $= -\infty$ (destroying capital) if $\mathrm{E}[\mu_t] - \mathrm{E}[\sigma_t^2/2] - \mathrm{E}[s_t] < 0$. ■

A.2 Proof of Theorem 3

Proof. We can rewrite (5) and (6) as

$$
d\begin{pmatrix} W_t \\ S_t \end{pmatrix} = A\begin{pmatrix} W_t \\ S_t \end{pmatrix}dt,
$$

where

$$
A = \begin{pmatrix} r & -1 \\ \kappa\tau & -\kappa \end{pmatrix}.
$$

This ODE can be solved by using an eigenvalue-eigenvector decomposition of A. The eigenvalues of A are the two roots of the eigenvalue equation $\det(A-\lambda I)=0$, given by

$$
\lambda = \frac{r-\kappa \pm \sqrt{(\kappa-r)^2 - 4\kappa(\tau-r)}}{2} = \frac{r-\kappa \pm \sqrt{(\kappa+r)^2 - 4\kappa\tau}}{2}.
$$

We will label the eigenvalues so that $\lambda_2 < 0 < \lambda_1$. The corresponding eigenvectors are given by $\phi_i = (1, r-\lambda_i)^\intercal$. The solution of the ODE is:

$$
\begin{pmatrix} W_t \\ S_t \end{pmatrix} = K_1 e^{\lambda_1 t}\phi_1 + K_2 e^{\lambda_2 t}\phi_2.
$$

The constants K_1 and K_2 can be determined by the initial conditions as $K_1 = (W_0\,(r-\lambda_2) - S_0)/(\lambda_1-\lambda_2)$ and $K_2 = (W_0\,(\lambda_1 - r) + S_0)/(\lambda_1-\lambda_2)$. Note that $0 < r-\lambda_1 < r-\lambda_2$, so that if $S_0/W_0 > r - \lambda_2$, then $K_2 > W_0$ and $K_1 = W_0 - K_2 < 0$, so wealth goes to zero in finite time and, thus, capital is not preserved in this case. Specifically, let the time that wealth reaches zero be t^*, then we have

$$W_t = K_1 e^{\lambda_1 t^*} + K_2 e^{\lambda_2 t^*} = 0 \Longleftrightarrow e^{(\lambda_1 - \lambda_2)t^*} = -\frac{K_2}{K_1} \Longleftrightarrow t^* = \frac{1}{\lambda_1 - \lambda_2} \log\left(-\frac{K_2}{K_1}\right),$$

where $\log(-K_2/K_1) > 0$ since $K_1 < 0 < K_2$ and $-K_1 = K_2 - W_0 < K_2$. ■

A.3 Proof of Theorem 4

We want to show that wealth in a unit of endowment hits zero in finite time with probability 1. We are studying the dynamics of spending and wealth given by (5) and (7) up until the first time when wealth hits zero, formally on the time interval

$$\mathcal{T} \equiv \begin{cases} [0, \inf\{t > 0 | w_t = 0\}] & \text{if } (\exists t > 0)(W_t = 0) \\ [0, \infty) & \text{otherwise} \end{cases}.$$

We need to show that with probability 1, $\mathcal{T}$ is a finite time interval.

The equations (5) and (7) taken together form a vector system of stochastic differential equations satisfying the usual Lipschitz conditions[12] that imply the system with initial conditions for wealth and spending at $t = 0$ have a unique continuous solution for all $t \in \mathcal{T}$. The spending dynamics (5) imply that we can write $S_t = S_0 e^{-\kappa t} + \kappa\tau \int_{s=0}^{t} W_s e^{-\kappa(t-s)} ds$, which is positive for all $t \in \mathcal{T}$. Therefore, the process $U_t \equiv W_t/S_t$ is well-defined and continuous for $t \in \mathcal{T}$. By Itô's lemma,

$$dU_t = \left(-1 + (\mu + \kappa)U_t - \kappa\tau U_t^2\right) dt + U_t \sigma dZ_t.$$

Recall that the first time U_t hits zero (if any) is the same as the first time W_T hits zero.

[12]Rogers and Williams (2000, Chapter V, Section 12) is more general than we need but handles the stopping at 0 nicely.

We use a martingale sample-path approach to proving our result; see Rogers and Williams (2000, IV.44-51).[13] Before providing the proof, we offer the following outline:

1. Find an increasing function F such that $F(U_t)$ is a local martingale (has zero drift). Any such F has $F(0)$ finite and $F(U)$ is unbounded as U increases. We also choose F such that $F(0) = 0$ and F is increasing. For our choice, F is a one-to-one continuous function mapping $[0, \infty)$ to $[0, \infty)$, so U_t hits zero when and only when $F(U_t)$ hits zero.

2. Find a change of time (an increasing continuous function rescaling the time variable) to convert $F(U_t)$ to have a variance of one per unit time, which implies that in the new time units the process equals the initial value plus a standard Wiener process. This is possible because $F(U_t)$ is a continuous local martingale.

3. For a standard Wiener process plus a positive initial value, we know it will hit zero in finite time, and we know a lot about the sample path up to this first hitting time. For example, we know the sample path is continuous and therefore almost surely bounded, and that the occupation time [14] between now and the hitting time is almost surely continuous. We also know the functional form of the expected occupation time.

4. The occupation time for $F(U_t)$ can be computed as a derivative of the change of time (which depends only on $F(U_t)$, not on t) times the occupation time for the time-changed process (the Wiener process). We can use this expression plus the known facts about the sample path of the Wiener process before hitting zero to prove $F(U_t)$ (and consequently U_t) hits zero in finite time.

[13] In essence, we are proving the "if" side of their Theorem V.51.2(ii) using the same approach but slightly different details. We cannot apply their result directly because we require a lot of notation and some preliminary results to use their result as stated.

[14] Occupation density is a version of the Radon-Nykodim derivative of the occupation measure, indicating the amount of time spent by stochastic process at a point/location during a period of time, e.g., $[0, t]$. It is also sometimes known as local time for real-valued stochastic processes.

To start with, we want to find a $\mathcal{C}^2$ function $F : \Re_{++} \to \Re$ such that $F(U_t)$ is a local martingale, i.e. has no drift. By Itô's lemma, we have

$$dF(U_t) = F'(U_t)\left[\left(-1+(\mu+\kappa)U_t-\kappa\tau U_t^2\right)dt+U_t\sigma dZ_t\right]+\frac{1}{2}F''(U_t)\left(\sigma U_t\right)^2 dt.$$

The drift of $F(U_t)$ is always 0 if and only if F satisfies

$$F'(u)\left(-1+(\mu+\kappa)u-\kappa\tau u^2\right)+\frac{1}{2}F''(u)\left(\sigma u\right)^2 = 0. \tag{16}$$

One solution is

$$F(U) = \int_{u=0}^{U} \exp\left(-\frac{2(\mu+\kappa)\log(u)}{\sigma^2}-\frac{2}{\sigma^2 u}+\frac{2\kappa\tau u}{\sigma^2}\right)du.$$

We will show momentarily that the integral exists, and given that existence, the condition (16) can be verified by direct calculation. For existence of the integral, first note that in the argument to the exponential, the term $-2/(\sigma^2 u)$ dominates when $u \downarrow 0$ (so the integrand tends to 0), and the term $2\kappa\tau u/\sigma^2$ dominates when u tends to infinity, so the integrand tends to infinity. Therefore, the integrand is finite, positive, and continuous everywhere, and the integral exists. Furthermore, since the integrand is always positive, $F'(u) > 0$ for $u > 0$, and since the integrand increases without bound as u increases, $\lim_{u\uparrow\infty} F(u) = \infty$. Furthermore, $F(0) = 0$.

Since $Q_t \equiv F(U_t)$ is a continuous local martingale, it is a time-changed Wiener process (perhaps on an augmented probability space). Specifically, there exist a Wiener process B_s starting at $B_0 = Q_0$ with variance one per unit time and a corresponding continuous and increasing time change $t = v(s)$ with $v(0) = 0$, such that $Q_{v(s)} = Q_0 + B_s$ up to the first time $Q_{v(s)}$ hits zero. Matching the cumulative variance, the time change can be computed implicitly as $\Sigma(v(s)) = s$, where the random process $\Sigma(t) = \int_{z=0}^{t} \mathrm{var}(dQ_z)$ is the increasing cumulative variance (quadratic variation) process for Q_t in the original time frame (which is not just linear since variance at t depends on U_t and is not constant). Applying Itô's lemma to $Q_t = F(U_t)$, we have

$$dQ_t = F'(U_t)\sigma U_t dZ_t \tag{17}$$

(the drift term cancels by construction) and therefore

$$\Sigma(t) = \int_{\tau=0}^{t} \left(F'(U_\tau)\right)^2 \sigma^2 U_\tau^2 d\tau = s. \tag{18}$$

Since F is one-to-one from $[0,\infty)$ onto itself, its inverse function $I(\cdot)$ is also one-to-one from $[0,\infty)$ onto itself. This allows a characterization of whether the boundary $U_t = 0$ is hit in finite time by checking whether Q_t hits zero in finite time.

In the time-changed version, B_s is a standard Wiener process, so B_s first hits zero at a random but finite time s, call it H_0. Although we know that H_0 is finite in the time-changed version, we need to show that this corresponds to a finite time in the original measure, which is not obvious since time may move more and more slowly in the original measure than in the new measure as B_s approaches 0. Note that in the original frame, (18) has the differential version

$$\left(F'(U_t)\right)^2 \sigma^2 U_t^2 dt = ds.$$

Now $U_t = I(B_s)$ and $t = v(s)$, so we can write

$$\frac{dt}{ds} = v'(s) = g(B_s), \tag{19}$$

where

$$g(q) \equiv \frac{1}{\left(F'(I(q))\right)^2 \sigma^2 I(q)^2}.$$

Therefore, the time until wealth hits zero in the original frame is

$$v(H_0) = \int_{s=0}^{H_0} v'(s)\, ds = \int_{s=0}^{H_0} g(B_s) ds, \tag{20}$$

and this is what we want to show to be finite. We prove this by changing the integral over time (20) into an integral over location, using the concept of local time (occupation density).

Let $l^q_{H_0}$ be the local time[15] of Wiener process B_s for any location q over the time interval $[0, H_0]$. Viewed as a function of q, the local time is the density of the measure of how much time the process spends at different values of q, which is what allows us to convert the integral over time (20) to an integral over space. Specifically,

$$v(H_0) = \int_{q=-\infty}^{\infty} l^q_{H_0} g(q)\, dq = \int_{q=0}^{Q_0} l^q_{H_0} g(q)\, dq + \int_{q=Q_0}^{\infty} l^q_{H_0} g(q)\, dq, \tag{21}$$

where we have used the fact that Q_t is positive on the time interval we are considering and therefore local time is zero for $q < 0$.

Consider first the second term on the rhs of (21). Note that B_s is a.s. continuous and therefore bounded on the compact interval $[0, H_0]$, and $l^q_{H_0}$ is continuous and equals zero outside of $[0, \max\{B_s | s \in [0, H_0]\}]$. Furthermore, it is easy to verify that $g(q)$ is a continuous function on any bounded interval of the form $[Q_0, q_{max}]$. Therefore, the second term is the integral of a continuous function on a compact interval and is therefore finite.

Now consider the first term on the rhs of (21). By Rogers and Williams (2000), V.51.1(i), for all $q > 0$, $\mathrm{E}[l^q_{H_0}] = 2\min(q, Q_0)$, which implies that the expectation of the first term is $2\int_{q=0}^{Q_0} qg(q)dq$. It suffices to show this expectation of the first term is finite, since that implies the first term is finite almost surely. Now, since $q = F(u)$,

$$\begin{aligned}
\int_{q=0}^{Q_0} qg(q)dq &= \int_{q=0}^{Q_0} \frac{1}{\left(F'(I(q))\sigma I(q)\right)^2} qdq \\
&= \int_{u=0}^{U_0} \frac{1}{(F'^2} F(u)F'(u)du \\
&= \int_{u=0}^{U_0} \frac{F(u)}{F'^2 u^2} du.
\end{aligned}$$

[15]The literature uses two different notions of local time that differ by a factor of two. We follow the convention in Rogers and Williams (2000).

All we have left to show that this integral is finite. The integrand is continuous on $(0, U_0]$, so it suffices to show that it has a finite limit at 0. Using L'Hôpital's rule and (16), we have

$$\begin{aligned}
\lim_{u\downarrow 0} \frac{F(u)}{F'^2 u^2} &= \lim_{u\downarrow 0} \frac{F'(u)}{F''^2 u^2 + 2F'^2 u} \\
&= \lim_{u\downarrow 0} \left(\sigma^2 u^2 F''(u)/F'^2 u\right)^{-1} \\
&= \lim_{u\downarrow 0} \left(\sigma^2 \left(-\frac{2(\mu+k)}{\sigma^2} u + \frac{2}{\sigma^2} + \frac{2\kappa\tau}{\sigma^2} u^2\right) + 2\sigma^2 u\right)^{-1} \\
&= \left(\sigma^2(-0 + \frac{2}{\sigma^2} + 0) + 0\right)^{-1} \\
&= \frac{1}{2}
\end{aligned}$$

which is finite. ■

A.4 Proof of Theorem 5

By Itô's lemma, the spending and wealth dynamics (8) and (7) with $\sigma_t = \sigma > 0$ constant imply that $\log(s_t) = \log(S_t/W_t)$ is an Ornstein-Uhlenbeck velocity process

$$d \log(s_t) = \kappa (\log \tau - \log(s_t))\, dt - \sigma dZ_t.$$

We will write s_t as a moving average for which the mean-square ergodic theorem applies. The moving average representation of the spending rate is

$$\log(s_t) = \log \tau - \sigma \int_{\Delta=-\infty}^{0} e^{\kappa\Delta} dZ_{t+\Delta}. \tag{22}$$

Hence, the process of $\log(s_t)$ is stationary with constant mean, variance, and autocovariance function

$$\begin{aligned}
\mathrm{E}\left[\log\left(s_t\right)\right] &= \log \tau, \\
\mathrm{var}\left[\log\left(s_t\right)\right] &= \frac{\sigma^2}{2\kappa}, \\
\mathrm{cov}\left[\log\left(s_t\right), \log\left(s_{t+\Delta}\right)\right] &= \frac{\sigma^2}{2\kappa} e^{-\kappa|\Delta|}.
\end{aligned}$$

As a result, s_t is lognormally distributed with mean, variance, and autocovariance function

$$\mathrm{E}\left[s_t\right] = \tau \exp\left(\frac{\sigma^2}{4\kappa}\right), \tag{23}$$

$$\mathrm{var}\left[s_t\right] = \tau^2 \left(\exp\left(\frac{\sigma^2}{2\kappa}\right) - 1\right) \exp\left(\frac{\sigma^2}{2\kappa}\right), \tag{24}$$

$$\mathrm{cov}\left[s_t, s_{t+\Delta}\right] = \tau^2 \left(\exp\left(\frac{\sigma^2}{2\kappa} e^{-\kappa|\Delta|}\right) - 1\right) \exp\left(\frac{\sigma^2}{2\kappa}\right). \tag{25}$$

Note that the autocovariance depends only on the lag $|\Delta|$ and not on time t. Therefore, s_t is also covariance stationary.

We now prove the spending rate s_t is a mean-square ergodic process as in Theorem 1. By the condition in the Theorem and (23), s_t is ergodic if the following limit is zero:

$$\begin{aligned}
\lim_{T\uparrow\infty} \frac{1}{T} \int_{\Delta=0}^{T} \left(\exp\left(\tfrac{\sigma^2}{2\kappa} e^{-\kappa|\Delta|}\right) - 1\right) d\Delta &= \lim_{T\uparrow\infty} \left(\frac{1}{T} \int_{\Delta=0}^{T^*} \left(\exp\left(\tfrac{\sigma^2}{2\kappa} e^{-\kappa|\Delta|}\right) - 1\right) d\Delta\right) \\
&\quad + \lim_{T\uparrow\infty} \left(\frac{1}{T} \int_{\Delta=T^*}^{T} \left(\exp\left(\tfrac{\sigma^2}{2\kappa} e^{-\kappa|\Delta|}\right) - 1\right) d\Delta\right) \\
&< 0 + \lim_{T\uparrow\infty} \left(\frac{1}{T}(T - T^*) \left(\exp\left(\tfrac{\sigma^2}{2\kappa} e^{-\kappa|T^*|}\right) - 1\right)\right) \\
&= \exp\left(\frac{\sigma^2}{2\kappa} e^{-\kappa|T^*|}\right) - 1,
\end{aligned}$$

where T^* is an arbitrary positive number, and the inequality follows because the integrand is decreasing in Δ. The integrand is positive, so the limit (if it exists) is nonnegative. Since T^* is arbitrary and the rhs decreases to 0 as T^* goes to infinity, the limit exists and must equal to zero, implying that s_t is an ergodic covariance-stationary process. The result follows from Theorem 2 (since μ_t and the constant σ^2 are also ergodic covariance-stationary processes) and (23). ■

A.5 Proof of Theorem 7

Let S_t^*, θ_t^*, and W_t^* be any feasible spending, investment in the risky asset, and wealth that may not satisfy (14). We want to match (or exceed) its value using the limit of strategies that do satisfy (14). Consider the alternative safe strategy

$$S_t^{\text{safe}} \equiv rW_0/2, \quad \theta_t^{\text{safe}} \equiv 0, \quad \text{and} \quad W_t^{\text{safe}} \equiv (1+e^{rt})W_0/2.$$

Then we will let

$$\begin{aligned} S_t^m &= (mS_t^* + S_t^{\text{safe}})/(m+1), \\ \theta_t^m &= (m\theta_t^* + \theta_t^{\text{safe}})/(m+1), \\ W_t^m &= (mW_t^* + W_t^{\text{safe}})/(m+1). \end{aligned}$$

It is easy to check that this is feasible and satisfies (14) for every m. Let M be the product of probability measure (across states) and Lebesgue measure (for positive times). Then, noting that probability measure integrates to one, we can write the expected utility of the safe strategy, S_t^{safe}, as

$$\int D_t u(S_t^{\text{safe}})dM = u(rW_0/2)\int_{t=0}^{\infty} D_t dt,$$

which is finite because $\int_{t=0}^{\infty} D_t dt$ and $u(rW_0/2)$ are both finite. In other words, $D_t u(S_t^{\text{safe}}) \in \mathcal{L}^1(M)$. Since the strategy (θ^*, S^*, W^*) has finite value, we also know that $D_t u(S_t^*) \in$

$\mathcal{L}^1(M)$. It also follows that

$$z_{min} \equiv \min(D_t u(S_t^{\text{safe}}), D_t u(S_t^*)) \in \mathcal{L}^1(M),$$

and

$$z_{max} \equiv \max(D_t u(S_t^{\text{safe}}), D_t u(S_t^*)) \in \mathcal{L}^1(M).$$

Since $(\forall m)(z_{min} \leq D_t u(S_t^m) \leq z_{max})$ and $D_t u(S_t^m)$ converges almost-surely to $D_t u(S_t^*)$, it follows that

$$\int D_t u(S_t^m) dM \to \int D_t u(S_t^*) dM, \text{ as } m \to \infty,$$

which is the required convergence of expected utility. ■