Portfolio Choice with Sustainable Spending: A Model of Reaching for Yield

John Y. Campbell and Roman Sigalov[1]

First draft: January 2020

[1]Campbell: Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, USA, and NBER. Email: john_campbell@harvard.edu. Sigalov: Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, USA. Email: rsigalov@g.harvard.edu. This paper originated in a May 2019 presentation by Campbell to the NBER Conference on Long-Term Asset Management, available online at https://scholar.harvard.edu/files/campbell/files/nber_ltamkeynoteslides.pdf. We are grateful to Ian Martin and Rob Sitkoff for helpful conversations on this topic.

Abstract

We show that reaching for yield—a tendency to take more risk when the real interest rate declines while the risk premium remains constant—results from imposing a sustainable spending constraint on an otherwise standard infinitely lived investor with power utility. This is true for two alternative versions of the constraint which make wealth and consumption follow martingales in levels or in logs, respectively. Reaching for yield intensifies when the interest rate is initially low, helping to explain the salience of the topic in the current low-rate environment. The sustainable spending constraint also affects the response of risktaking to a change in the risk premium, which can even be negative when the riskless interest rate is sufficiently low.

1 Introduction

How does the level of the safe real interest rate affect investors' willingness to take risk? The conventional answer in finance theory, derived by Merton (1969, 1971), is that it does not. A long-lived investor with constant relative risk averse power utility, facing a constant risk premium, allocates a constant share of wealth to a risky asset regardless of the level of the real interest rate.[2]

Contrary to this conventional theory, the prolonged period of low interest rates in the last decade has led some observers to believe that investors "reach for yield", taking more risk when the real interest rate declines. Reaching for yield has been a particular concern among central bankers who fear that risk-taking may be an unintended consequence of loose monetary policy (Rajan 2006, 2013, Borio and Zhu 2012, Stein 2013).

Most of the literature on reaching for yield studies the responses to interest rates of financial intermediaries such as banks, money market funds, and pension funds (Maddaloni and Peydró 2011, Chodorow-Reich 2014, Jiménez et al 2014, Hanson and Stein 2015, Andonov, Bauer, and Cremers 2017, Di Maggio and Kacperczyk 2017, Lu et al 2019).[3] To the extent that models of reaching for yield are offered, they rely on institutional frictions specific to these intermediaries such as bank liquidity management (Drechsler, Savov, and Schnabl 2018, Acharya and Naqvi 2019), a zero lower bound on the nominal return that banks and money funds can offer depositors (Chodorow-Reich 2014, Di Maggio and Kacperczyk 2017), or the discounting of pension fund liabilities when calculating funding status (Andonov, Bauer, and Cremers 2017, Lu et al 2019). An exception to this focus on intermediaries is Lian, Ma, and Wang (2019) who document reaching for yield behavior among households and suggest a behavioral explanation.

[2]Merton's analysis, like this paper, is set in continuous time. Campbell (2018) provides a textbook exposition of the result in discrete time and shows how it extends from power utility to the related model where an investor has Epstein-Zin preferences.

[3]There are also papers that use the term "reaching for yield" in a different sense, to refer to an unconditional propensity for financial intermediaries to take risk (Becker and Ivashina 2015, Choi and Kronlund 2017). In this paper we use the term exclusively to refer to the response of risk-taking to the interest rate.

In this paper we present a simple theoretical model that generates reaching for yield through a different mechanism. Our model applies most naturally to endowments and trusts, a different type of institutional investor than has been considered in the existing literature.

We start from the classic Merton model of an infinitely lived investor with power utility who chooses consumption and the portfolio allocation to a risky asset and a safe asset, both of which have iid returns. To this classic model we add one constraint, that the investor must consume in each period the expected return on the portfolio. This level of spending ensures that both consumption and wealth follow martingales, so the investor cannot plan either to run down wealth or to accumulate it. For this reason we call the constraint a "sustainable spending" constraint.

One natural interpretation of the sustainable spending constraint is that it reflects a promise that institutions with endowments, such as universities, make to donors. A gift to the endowment implies a commitment by the university to undertake an activity not just for a few years, but permanently. Donors who wish to have a greater short-term impact can make current-use gifts, but endowment gifts reflect a donor's desire for permanent impact.[4] Spending more than the expected return would run down the endowment and support only temporary spending, whereas spending less than the expected return would imply that the gift has less immediate impact than is needed for sustainability. In addition, spending less than the expected return would cause the endowment to grow on average, which can lead to unwelcome attention and even unfavorable tax consequences in the long run.[5]

The sustainable spending constraint is also a way to model the legal restrictions on endow-

[4] As Harvard University puts it, "The University's spending practice has to balance two competing goals: the need to fund the operating budget with a stable and predictable distribution, and the obligation to maintain the long-term value of endowment assets after accounting for inflation." (https://www.harvard.edu/about-harvard/harvard-glance/endowment, accessed January 28, 2020.) A literature on endowment spending has discussed smoothing rules that stabilize the operating budget but has paid less attention to the long-run spending rate. See for example Dybvig (1995), Brown et al (2014), and Gilbert and Hrdlicka (2015). Here we focus on the determination of the long-run spending rate and its interaction with asset allocation.

[5] The tax reform passed by Congress in 2017 includes an excise tax on private universities with over 500 tuition-paying students and assets of over $500,000 per student.

ments. Under the Uniform Prudent Management of Institutional Funds Act (UPMIFA), which governs endowment expenditures for nonprofit and charitable organizations in all US states except Pennsylvania, endowment spending must be consistent with the long-run preservation of capital but need not keep the endowment in any particular relation to the original historic value of gifts as was required under earlier state laws. Similar restrictions have been imposed on some sovereign wealth funds such as the Norwegian oil fund.[6]

The sustainable spending constraint may also describe the legal restrictions on trustees in cases where the principal of a trust is reserved for one beneficiary but income is paid to another. Under the 1997 revision of the Uniform Principal and Income Act, trustees are no longer bound to distribute only dividends or coupon payments to income beneficiaries, but have the "power to adjust" income distributions to reflect the expected return on investments (Sitkoff and Dukeminier 2017, pp. 669–670). If a trustee invests in the interests of a long-lived income beneficiary (that is, if the date of principal transfer is in the distant future), then the trustee's decisions may be well approximated by the model developed here.

Finally, the sustainable spending constraint might be interpreted more broadly as a parsimonious way to describe the behavior of older households who live off financial wealth but fear the consequences of running down that wealth, for example because they wish to preserve funds to cover unexpected medical expenses or to leave a bequest. Such households may be willing to take investment risk but may be unwilling to adopt a consumption plan that implies expected declines in future wealth.

[6]Norges Bank Investment Management, the manager of Norway's Government Pension Fund Global (the official name of the oil fund) explains the restrictions in these words: "So that the fund benefits as many people as possible in the future too, politicians have agreed on a fiscal rule which ensures that we do not spend more than the expected return on the fund. On average, the government is to spend only the equivalent of the real return on the fund, which is estimated to be around 3% per year. In this way, oil revenue is phased only gradually into the economy. At the same time, only the return on the fund is spent, and not the fund's capital." (https://www.nbim.no/en/the-fund/about-the-fund/, accessed January 26, 2020.) The 3% figure is a current number; 4% was assumed until the spring of 2017.

The sustainable spending constraint breaks the separation between consumption/saving and asset allocation decisions that underlies the classic Merton result. When the interest rate declines, an investor with a constant rate of time preference wants to increase the expected growth rate of marginal utility. In the classic model, the investor achieves this by consuming more today and less in the future, but the sustainable spending constraint makes this impossible. Instead, the sustainable-spending constrained investor reaches for yield, taking more risk as a way to increase consumption today while paying the cost—volatile future consumption—only in the future.

We find that reaching for yield intensifies when the interest rate is initially low. This may help to explain why the phenomenon is more widely discussed today than it was in the higher-rate environment of the late 20th Century. The sustainable spending constraint also affects the response of risktaking to a change in the risk premium. This response declines as the interest rate declines and even becomes negative at sufficiently low levels of interest rates.

A subtle issue that we discuss in this paper is whether the expected return that governs sustainable spending is the expected simple (arithmetic average) return or the expected log (geometric average) return. In the former case, the levels of wealth and consumption are martingales while in the latter case the logs of wealth and consumption are martingales. The former case may seem more natural at first, but it implies that wealth and consumption approach zero with ever-increasing probability over time. As Dybvig and Qin (2019) emphasize, this is inconsistent with the spirit of the sustainable-spending constraint. Therefore, while we explore both cases in this paper, we emphasize the results for the geometric average case.

The organization of the paper is as follows. In section 2 we set up the problem and show how the sustainable-spending constraint forces the investor to trade off current (and expected future) consumption against the volatility of consumption. We describe the solution using indifference curves and a feasible curve derived from portfolio choice. In section 3 we explore the arithmetic-average model in greater detail, presenting a closed-form solution for the portfolio rule with a sustainable-spending constraint. In section 4 we analyze the geometric-average model, proving that the portfolio rule displays reaching for yield even though we do not have a closed-form solution.

Sections 2–4 present comparative static results, comparing economies with different constant levels of interest rates. In section 5, by contrast, we consider a dynamic model where the riskless interest rate follows a persistent stochastic process while the risk premium remains constant. In this model the portfolio rule depends on the correlation between the risky return and the innovation to the riskless interest rate, as in Merton (1973), Campbell and Viceira (2001), and other models of intertemporal hedging demand for risky assets. Section 6 concludes, and an online appendix (Campbell and Sigalov 2020) contains additional details of the analysis.

2 Trading Off the Level and Volatility of Consumption

2.1 Preferences and Indifference Curves

We start our analysis by outlining the tradeoff faced by an investor who follows a sustainable consumption rule. We consider an investor with time-separable utility of consumption $u(c_t)$ and time discount rate ρ whose log consumption follows a Brownian motion with an arbitrary drift μ_c and constant volatility σ_c:

$$d \log c_t = \mu_c dt + \sigma_c dZ_t. \tag{1}$$

A straightforward application of Ito's lemma tells us the stochastic process followed by the level of consumption:

$$dc_t = c_t \left(\mu_c + \frac{1}{2} \sigma_c^2 \right) dt + c_t \sigma_c dZ_t. \tag{2}$$

Using this law of motion for consumption, we can derive the Hamilton-Jacobi-Bellman (HJB) equation describing the value function $v(c_t)$ for an infinitely lived investor with time-separable utility and time discount rate ρ:

$$\rho v(c_t) = u(c_t) + v'(c_t)c_t\left(\mu_c + \frac{1}{2}\sigma_c^2\right) + \frac{1}{2}v''(c_t)^2\sigma_c^2. \tag{3}$$

We now assume that the investor has instantaneous power utility of consumption with relative risk aversion γ. Under this assumption the natural guess for the form of the value function is

$$v(c_t) = \theta\frac{c_t^{1-\gamma}}{1-\gamma}, \tag{4}$$

for some parameter θ. Substituting equation (4) into equation (3) and simplifying, we can solve for θ as

$$\theta = \left(\rho + (\gamma-1)\mu_c - (\gamma-1)^2\frac{\sigma_c^2}{2}\right)^{-1}. \tag{5}$$

We will consider the empirically relevant case where the investor's coefficient of relative risk aversion $\gamma > 1$. In this case θ in (5) decreases with μ_c and increases with σ_c; but it is multiplied by a negative number in (4) to deliver a negative value function that, as one would expect, increases with μ_c and decreases with σ_c.

An important implication of equation (5) is that there is an upper bound on the volatility of consumption that is consistent with finite utility of the infinitely lived investor. When σ_c exceeds this upper bound, θ becomes negative and the value function is undefined.

Putting these results together, we have shown that the value function for a power-utility investor with an exogenous stream of consumption given by equation (1) is

$$v(c_0) = E_0\int_0^\infty e^{-\rho t}u(c_t)dt = \left(\rho + (\gamma-1)\mu_c - (\gamma-1)^2\frac{\sigma_c^2}{2}\right)^{-1}\frac{c_0^{1-\gamma}}{1-\gamma}. \tag{6}$$

If we fix the initial level of the value function at some constant v, then equation (6) can be rewritten as an indifference condition relating the initial level of consumption c_0 and the volatility of log consumption σ_c:

$$c_0 = \left[\left(\rho + (\gamma - 1)\mu_c - (\gamma - 1)^2 \frac{\sigma_c^2}{2} \right) (1 - \gamma) v \right]^{\frac{1}{1-\gamma}}. \tag{7}$$

This equation tells us that c_0 is increasing in both σ_c and v: a larger initial level of consumption is required to compensate the investor for higher volatility given constant value, or to deliver higher value with constant volatility. Hence, if we plot indifference curves in (σ_c, c_0) space, each curve is upward sloping and the agent will pick the highest possible curve subject to a constraint that we next derive from the portfolio choice problem.

2.2 Sustainable Spending Rules

Sustainable spending rules determine the drift in log consumption, μ_c. They also impose constraints on the relationship between c_0 and σ_c, since an increase in consumption today can be financed only by taking more portfolio risk which implies greater volatility in consumption growth.

Arithmetic sustainable spending rule An arithmetic sustainable spending rule implies that the drift in the level of consumption is zero: $\mu_c + \sigma_c^2/2 = 0$ in equation (2). In this case the indifference condition (7) becomes

$$c_0 = \left[\left(\rho - \gamma(\gamma - 1) \frac{\sigma_c^2}{2} \right) (1 - \gamma) v \right]^{\frac{1}{1-\gamma}}. \tag{8}$$

Normalizing initial wealth $w_0 = 1$, the arithmetic rule implies that the relationship between the initial level of consumption and the volatility of consumption is

$$c_0 = r_f + \left(\frac{\mu}{\sigma}\right)\sigma_c, \tag{9}$$

where r_f is the riskless interest rate, μ is the expected excess return on a risky asset, and σ is the standard deviation of the risky asset return. Equation (9) follows from the familiar relationships in a Merton model that $\sigma_c = \alpha\sigma$ and the arithmetic expected portfolio return is $r_f + \alpha\mu$, where α is the portfolio share in a risky asset. It implies a linear tradeoff between the initial level and the volatility of consumption with intercept r_f and slope (μ/σ).

Although the arithmetic rule is simple and intuitive, and its linear constraint is analogous to the one used in classic mean-variance analysis when a riskless asset is available, it has an important disadvantage emphasized by Dybvig and Qin (2019). The arithmetic rule implies that the drift in log consumption is negative: $\mu_c = -\sigma_c^2/2$. The solution to equation (1) is then

$$c_t = c_0 \exp\left\{-\frac{1}{2}\sigma_c^2 t + \sigma_c Z_t\right\}. \tag{10}$$

When we take the limit of this expression as $t \to \infty$, using the fact that $Z_t/t \to 0$ almost surely, we conclude that $c_t \to 0$ almost surely. Such a property is undesirable since it implies that spending will eventually approach zero as the investor exhausts available wealth. Intuitively, the right-skewed distribution of the level of consumption, with increasing variance as the horizon increases, implies that almost all probability mass approaches zero while the constant expectation of future consumption is sustained by a vanishingly small number of scenarios in which consumption is extremely high.[7]

[7] An alternative intuition is that with an arithmetic sustainable spending rule the level of consumption is a martingale bounded below by zero. Like any bounded martingale, it must converge almost surely, but the only level to which it can converge is zero because this is the only level at which the consumption process has zero volatility. Martin (2012) presents a related analysis.

Geometric sustainable spending rule A geometric sustainable spending rule sets the drift in log consumption equal to zero: $\mu_c = 0$ in equation (1). This implies that log consumption is a martingale and, since log consumption is conditionally normally distributed at all horizons, the median value of future consumption equals the current value of consumption. This behavior seems more in accord with the spirit of a sustainable spending constraint.

Under the geometric sustainable spending rule, the indifference condition (7) becomes

$$c_0 = \left[\left(\rho - (\gamma - 1)^2 \frac{\sigma_c^2}{2} \right) (1 - \gamma) v \right]^{\frac{1}{1-\gamma}}. \tag{11}$$

Comparing the geometric indifference condition (11) with the arithmetic indifference condition (8), the difference is in the coefficient multiplying $\sigma_c^2/2$. Since $\gamma(\gamma - 1) > (\gamma - 1)^2$ under our maintained assumption that $\gamma > 1$, for any given value v the indifference curve for the arithmetic rule lies above the indifference curve for the geometric rule, except when $\sigma_c = 0$ where the two curves converge.

Normalizing initial wealth $w_0 = 1$, the geometric rule implies that the relationship between the initial level of consumption and the volatility of consumption is

$$c_0 = r_f + \left(\frac{\mu}{\sigma} \right) \sigma_c - \frac{1}{2} \sigma_c^2, \tag{12}$$

where the last term is the Jensen's Inequality difference between the arithmetic and geometric mean of consumption growth. Equation (12) is a concave rather than a linear constraint. As volatility increases, it has a diminishing effect on the expected log portfolio return and therefore on the initial level of consumption.

The highest level of initial consumption is obtained when the investor holds the growth-optimal portfolio with maximum log return. In this case $\sigma_c = \mu/\sigma$ and

$$c_0 = r_f + \frac{1}{2} \left(\frac{\mu}{\sigma} \right)^2. \tag{13}$$

As in the standard Merton problem, we can show that as γ approaches one, the optimal portfolio approaches a growth optimal portfolio and consumption will be given by equation (13).

2.3 Graphical Analysis

To understand the properties of indifference curves with sustainable spending constraints, we can differentiate equations (8) or (11) with respect to the standard deviation of consumption, σ_c, to see that under either type of constraint, the slope of the indifference curve is positive and increasing in the standard deviation σ_c. The slope is zero—the indifference curve is flat—at zero standard deviation, reflecting the fact that investors with twice differentiable utility always take some amount of any compensated risk. The slope approaches infinity as the standard deviation approaches the upper bound at which the value function no longer converges. Since as already noted $\gamma(\gamma-1) > (\gamma-1)^2$ under our maintained assumption that $\gamma > 1$, the slope is everywhere greater and the upper bound is smaller for the arithmetic constraint than for the geometric constraint.

These properties are illustrated in Figure 1. There, we plot indifference curves and spending constraints for both types of constraint and a common level of value v.[8] The two indifference curves have the same intercept with the vertical axis, but the arithmetic indifference curve is higher and steeper elsewhere and its asymptote lies to the left of that for the geometric indifference curve.

Figure 1 also illustrates the fact that at the optimum for the geometric rule, that is for the value v where the geometric indifference curve is tangent to its concave portfolio constraint, the arithmetic indifference curve lies above its linear portfolio constraint. To obtain the optimal choice under the arithmetic constraint we need to shift the indifference curve down, corresponding to a lower value, in order to satisfy the portfolio constraint. In other words optimized consumption delivers lower lifetime utility under the arithmetic constraint than under the geometric constraint;

[8] The parameters assumed are a 2% riskfree interest rate, a 6% risk premium, an 18% standard deviation of the risky return (implying a Sharpe ratio of 1/3), a 5% rate of time preference, and risk aversion of 2.

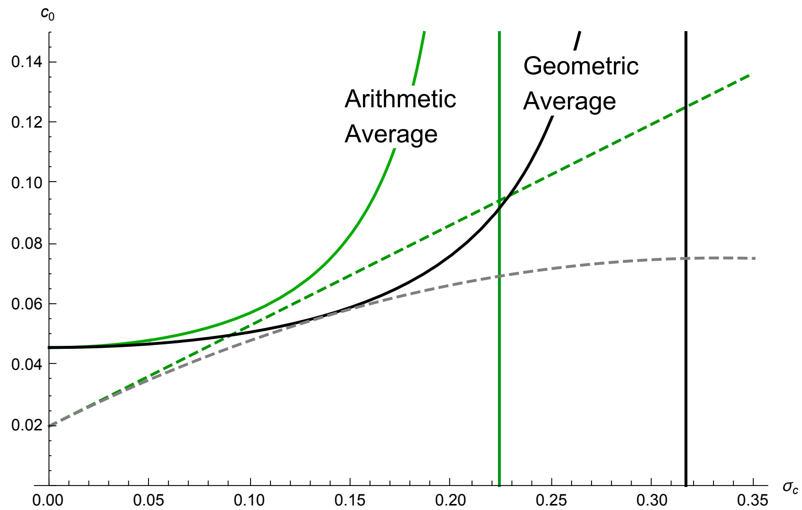

Figure 1: Comparing Arithmetic and Geometric Average Spending Constraints

this may be an additional reason to prefer the geometric formulation of the sustainable spending constraint.

Existence of a solution Figure 1 allows us to understand the conditions for the existence of a solution: a point where an indifference curve is tangent to the spending constraint. Since all indifference curves must be in the positive quadrant, a necessary condition for a tangency point to exist is that the spending constraint has a maximum above zero. This condition is trivially satisfied by the arithmetic constraint, which is linear with a positive slope, but not by the concave geometric constraint. If the riskfree rate is low enough, then even when the investor holds the growth-optimal portfolio, the expected portfolio return may be negative and no sustainable solution will exist for a geometric spending constraint. This problem arises if

$$r_f < -\frac{1}{2}\left(\frac{\mu}{\sigma}\right)^2. \tag{14}$$

A solution may also fail to exist if the spending constraint becomes positive only at levels of consumption volatility σ_c that exceed the upper bound for the value function to be finite. This problem can arise either for the arithmetic constraint or for the geometric constraint. In the geometric case, from equation (11) the highest permissible level of consumption volatility is

$$\sigma_c^\star = \sqrt{\frac{2\rho}{(1-\gamma)^2}}, \tag{15}$$

and from equation (12) the portfolio constraint intersects the x-axis at the point

$$\sigma_c^{\star\star} = \mu/\sigma - \sqrt{\mu^2/\sigma^2 + 2r_f}. \tag{16}$$

If $\sigma_c^{\star\star} > \sigma_c^\star$, the solution will not exist.

Reaching for yield The graphical analysis can be used to understand why sustainable spending rules generate reaching for yield. A change in the riskfree interest rate, with no change in the risk premium, is a parallel shift up or down in the spending constraint. An increase in the riskfree rate is an upward shift that improves the opportunity set and increases the achievable value v, while a decrease is a downward shift that reduces v.

Equations (8) and (11) imply that for both the arithmetic and geometric cases, the slope of the indifference curve is increasing in v. Mathematically, this is because v enters multiplicatively in these equations, increasing both the value and the slope of the indifference curve for any value of consumption volatility.

The response of the slope of the indifference curve to the level of value v implies that as the spending constraint shifts up, the optimal standard deviation of consumption declines; while as the spending constraint shifts down, the optimal standard deviation of consumption increases. This is precisely reaching for yield. It is illustrated in Figure 2 for the geometric case.[9]

[9] All parameters are the same as in Figure 1 except that we consider riskfree rates of 2.5%, 0%, and −2.5%.

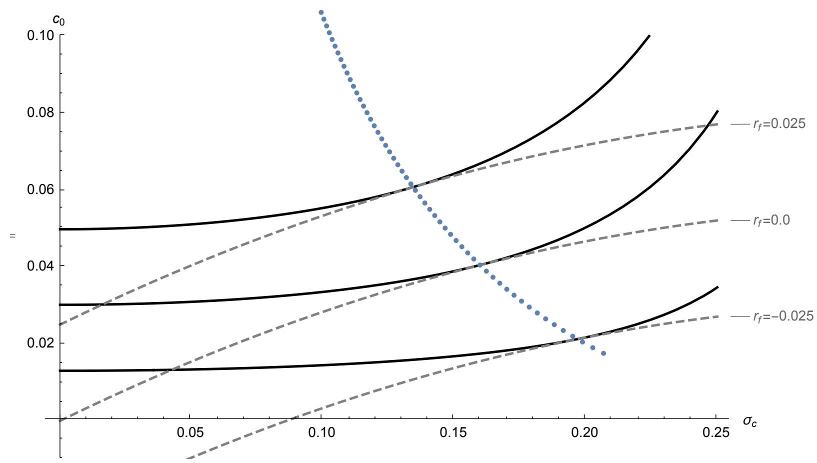

Figure 2: Reaching for Yield. The Effect of the Riskfree Rate on Initial Consumption and Consumption Volatility

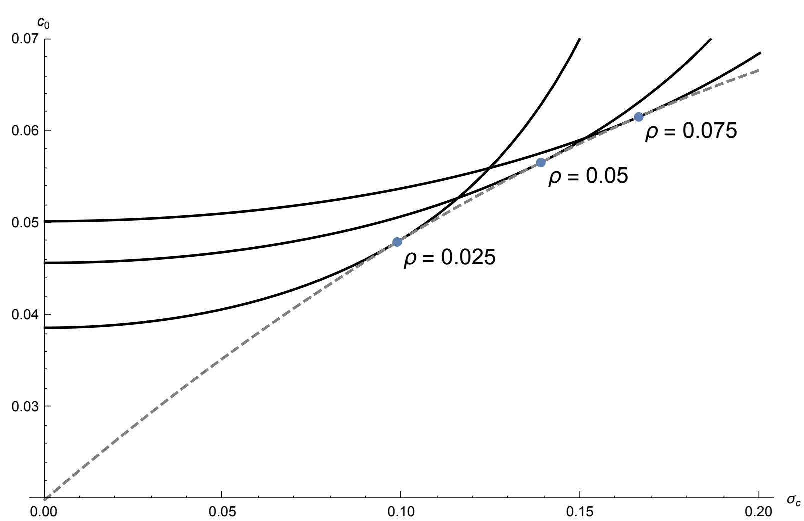

Figure 3: The Effect of Impatience on Initial Consumption and Consumption Volatility

Reaching for yield reflects the fact that in the problem with a sustainable spending constraint both consumption today and the stability of the consumption path are normal goods, so the investor uses an improvement in investment opportunities to increase both of them, in other words to reduce volatility as well as to increase current consumption.

The effect of impatience We can also analyze the effect of impatience on the problem. Equations (8) and (11) imply that for both the arithmetic and geometric cases, the slope of the indifference curve is decreasing in the time discount rate ρ. A more impatient investor values current consumption more relative to volatility that is realized in the future. Hence, the optimal solution will involve higher current consumption and higher consumption volatility as illustrated in Figure 3 for the geometric case.[10]

3 Arithmetic Average Model

Having presented a preliminary graphical analysis of the problem, we now write it as a standard portfolio choice problem. This section imposes an arithmetic average sustainable consumption constraint, while the next section uses a geometric average constraint.

The arithmetic constraint requires that the flow of consumption equals the investor's wealth times the simple expected return on the agent's portfolio:

$$c_t dt = w_t \mathrm{E} dr_{p,t}. \tag{17}$$

Here

$$dr_{p,t} = \alpha dr_t + (1-\alpha) r_f dt, \tag{18}$$

[10] All parameters are the same as in Figure 1 except that we consider time preference rates of 2.5%, 5%, and 7.5%.

where α is the portfolio share in a risky asset whose return dr_t is

$$dr_t = (r_f + \mu)dt + \sigma dZ_t. \tag{19}$$

As discussed earlier r_f is the riskfree interest rate, μ is the constant risk premium, and σ is the volatility of the risky asset return. The Sharpe ratio of the risky asset is the ratio μ/σ.

Combining these equations we have that

$$c_t dt = w_t(r_f + \alpha\mu)dt. \tag{20}$$

If we substitute this consumption rule into the budget constraint

$$dw_t = w_t dr_{p,t} - c_t dt, \tag{21}$$

we have

$$dw_t = w_t \alpha\sigma dZ_t, \tag{22}$$

showing that wealth follows a martingale with volatility proportional to the level of wealth. Since consumption is a constant fraction of wealth, it too follows a martingale.

Under the arithmetic constraint the maximization problem of the agent is

$$\max_{\alpha_t} \int_0^\infty e^{-\rho t} \frac{c_t^{1-\gamma}}{1-\gamma} dt \ = \int_0^\infty e^{-\rho t} \frac{\left(w_t(r_f + \alpha\mu)\right)^{1-\gamma}}{1-\gamma} dt \tag{23}$$

subject to equation (22).

We guess that the value function takes the form

$$v(w_t) = A \frac{w_t^{1-\gamma}}{1-\gamma}, \tag{24}$$

where A is a constant representing investment opportunities that depends on the riskfree rate and the Sharpe ratio of the risky asset. Equation (24) implies that the HJB equation takes the form

$$\begin{aligned} \rho v(w_t) &= \max_{\alpha} \left\{ \frac{(w_t(r_f + \alpha\mu))^{1-\gamma}}{1-\gamma} + \frac{1}{2}(-\gamma A w_t^{-\gamma-1}) w_t^2 \alpha^2 \sigma^2 \right\} \\ &= w_t^{1-\gamma} \max_{\alpha} \left\{ \frac{(r_f + \alpha\mu)^{1-\gamma}}{1-\gamma} - \frac{1}{2}\gamma A \alpha^2 \sigma^2 \right\}. \end{aligned} \tag{25}$$

The corresponding first-order condition is

$$(r_f + \alpha\mu)^{-\gamma}\mu - A\sigma^2\alpha = 0. \tag{26}$$

Closed-form solution Using equations (25) and (26) we can solve for the risky share α explicitly as

$$\alpha = \frac{-r_f + \sqrt{\mathcal{D}}}{\mu(1+\gamma)}, \tag{27}$$

where

$$\mathcal{D} = r_f^2 + 2\rho \left(\frac{1+\gamma}{\gamma}\right) \left(\frac{\mu}{\sigma}\right)^2. \tag{28}$$

This solution has several standard properties. Portfolio volatility, $\alpha\sigma = \sigma_c$, depends on the risk premium and volatility of the risky asset only through the Sharpe ratio (μ/σ). This is because leverage that alters the volatility of the risky asset without changing its Sharpe ratio does not affect the investor's opportunity set and hence does not alter the volatility of the investor's portfolio or consumption. Also, the risky share α declines with the volatility σ of the risky asset and with risk aversion γ when other parameters of the model are fixed (although it is not inversely proportional to σ^2 or γ as would be the case in the standard model).

The solution also has several nonstandard properties summarized in the following proposition.

Proposition 1 *(Comparative Statics for the Arithmetic Average Model). In the arithmetic average model, the risky share α has the following properties.*

1. *α is a decreasing and convex function of the riskfree rate r_f.*
2. *α is an increasing function of the rate of time preference ρ.*
3. *α is an increasing function of the risk premium μ when $r_f > 0$, and a decreasing function of μ when $r_f < 0$.*

Proof: See the online appendix, Campbell and Sigalov (2020).

The first property says that the investor reaches for yield, taking more risk when the riskfree rate is low. Furthermore, the risky share is convex in the riskfree rate, implying that reaching for yield is stronger when the interest rate is low. This result may help to explain why reaching for yield is so actively discussed in today's low-rate environment, and why this paper has been written in 2020 even though the technology to write it was already standard decades ago.

The second property says that a more impatient investor takes more risk. Intuitively, risktaking increases current consumption and the volatility that this implies for future consumption is heavily discounted by an impatient investor.

The third property says that the riskfree rate influences the response of risktaking to the risk premium, which has the same sign as the riskfree rate. This response is constant and positive in the standard model, but once a sustainable spending constraint is imposed an increase in the risk premium has offsetting income and substitution effects on risktaking. The substitution effect is to steepen the portfolio constraint, which by itself would increase risktaking. The income effect is to raise the portfolio constraint, which reduces risktaking in the same way as an increase in the riskless interest rate. The income effect becomes more powerful at low levels of the riskfree rate, both because the upward shift of the portfolio constraint is larger at a higher initial level of portfolio risk, and because the reaching for yield effect of a given upward shift is stronger when the riskfree rate is low.

The two effects of the risk premium offset one another when the riskfree interest rate is zero. With $r_f = 0$ the solution (27) becomes

$$\alpha = \frac{1}{\sigma}\sqrt{\frac{2\rho}{\gamma(1+\gamma)}} \tag{29}$$

for any nonzero risk premium μ, which depends positively on the rate of time preference and negatively on risk and risk aversion, but not on the level of μ.[11]

4 Geometric Average Model

We now consider the geometric average portfolio constraint. The value of the investor's portfolio, without subtracting consumption, follows the process

$$\frac{dV_t}{V_t} \equiv (r_f + \alpha\mu)dt + \alpha\sigma dZ_t. \tag{30}$$

The geometric sustainable consumption rule is

$$\begin{aligned} c_t dt &= w_t \mathrm{E}[d\log V_t] \\ &= w_t \left[r_f + \alpha\mu - \frac{1}{2}\alpha^2\sigma^2 \right] dt, \end{aligned} \tag{31}$$

where the second line is obtained by applying Ito's lemma to equation (30).

[11] With $r_f = 0$ and $\mu = 0$ the problem has no solution because no positive level of consumption is sustainable.

If we substitute this consumption rule into the budget constraint (21) and simplify we obtain

$$dw_t = w_t \frac{1}{2}\alpha^2\sigma^2 dt + w_t \alpha\sigma dZ_t, \tag{32}$$

and using Ito's lemma once again we get

$$d\log(w_t) = \alpha\sigma dZ_t. \tag{33}$$

This shows that in the geometric average model, the log and not the level of wealth follows a random walk. Since log consumption equals log wealth plus a constant, it too follows a random walk.

Under the geometric constraint the maximization problem of the agent is

$$\max_{\alpha_t} \int_0^\infty e^{-\rho t} \frac{c_t^{1-\gamma}}{1-\gamma} dt \;=\; \int_0^\infty e^{-\rho t} \frac{\left(w_t[r_f + \alpha\mu - \frac{1}{2}\alpha^2\sigma^2]\right)^{1-\gamma}}{1-\gamma} dt \tag{34}$$

subject to equation (32).

We guess that the value function takes the same form as in the arithmetic model, equation (24). This implies that the HJB equation for the geometric model is

$$\rho v(w_t) = w_t^{1-\gamma} \max_{\alpha} \left\{ \frac{(r_f + \alpha\mu - \frac{1}{2}\alpha^2\sigma^2)^{1-\gamma}}{1-\gamma} + \frac{1}{2}(1-\gamma)A\alpha^2\sigma^2 \right\}, \tag{35}$$

with first-order condition

$$\left(r_f + \alpha\mu - \frac{1}{2}\alpha^2\sigma^2 \right)^{-\gamma} + (1-\gamma)A\alpha\sigma^2 = 0. \tag{36}$$

Unlike the arithmetic average model, the geometric average model does not admit a closed form solution. However, by combining equations (35) and (36) and using the implicit function theorem for two equations and the model's second-order condition we can prove the following proposition.

Proposition 2 *(Comparative Statics for the Geometric Average Model). In the geometric average model, the risky share α has the following properties.*

1. *α is a decreasing function of the riskfree rate r_f.*
2. *α is an increasing function of the rate of time preference ρ.*
3. *There exists $r_f^* < 0$ such that for $r_f > r_f^*$, α is an increasing function of the risk premium μ and for $r_f < r_f^*$, α is a decreasing function of μ.*
4. *As risk aversion γ approaches 1, α approaches the growth optimal level μ/σ^2 for all values of r_f.*

Proof: See the online appendix, Campbell and Sigalov (2020).

The first property is a weaker version of the first property in Proposition 1. We have not been able to prove that α is convex in r_f, although this does seem to be the case in numerical examples. The second property is the same as that in Proposition 1. The third property is similar to the third property in Proposition 1, but in the geometric case the interest rate that causes a sign switch in the effect of the risk premium on risktaking is not necessarily zero as it is in the arithmetic case. The fourth property is new to the geometric case, and reflects the fact that with a geometric constraint the growth-optimal portfolio (which would be optimal in the standard model) also maximizes current consumption. We now discuss these properties in greater detail.

Reaching for yield In the standard portfolio choice problem the optimal risky share is $\alpha = \mu/\gamma\sigma^2$ resulting in an expected (arithmetic average) portfolio return of

$$\mathrm{E}r_p = r_f + \alpha\mu = r_f + \frac{\mu^2}{\gamma\sigma^2}. \tag{37}$$

This implies that the derivative of the expected portfolio return with respect to the riskless interest rate is one. If we plot the expected portfolio return $\mathrm{E}r_p$ against r_f we get a 45-degree line with a positive intercept.

In the portfolio choice problem with a geometric sustainable spending constraint, a lower riskfree rate leads to a higher risky share. This implies that so that

$$\frac{\partial \mathrm{E}r}{\partial r_f} = 1 + \underbrace{\frac{\partial \alpha(r_f)}{\partial r_f}}_{<0} \mu < 1. \tag{38}$$

If we plot $\mathrm{E}r_p$ against r_f for the geometric problem, the resulting curve will be flatter than a 45-degree line.

We illustrate this property in Figure 4 where we plot the expected portfolio return against the riskfree rate r_f for several different portfolios.[12] The figure includes a 45-degree line going through the origin, representing a portfolio with $\alpha = 0$ that is fully invested in the riskfree asset. It also includes a 45-degree line with a positive intercept of μ^2/σ^2 corresponding to the growth-optimal portfolio that will be held by an investor with log utility. A conservative investor with $\gamma > 1$ and a geometric sustainable spending constraint picks a portfolio whose expected return lies between these two lines.

Figure 4 illustrates the optimally chosen expected return for investors with $\gamma = 2$ (upper orange curve) and $\gamma = 3$ (lower blue curve). The slope of these curves is less than one, and for a more aggressive investor with $\gamma = 2$ the slope actually becomes negative when the riskfree rate becomes sufficiently low. In this region of the parameter space the tendency to reach for yield is so strong that a small decrease in the risk free rate can actually increase the arithmetic expected portfolio return (although it does not increase the geometric expected portfolio return and therefore does not increase the investor's current consumption).

[12] All parameters are the same as in Figure 1 except that we consider risk aversion levels of 1, 2, and 3.

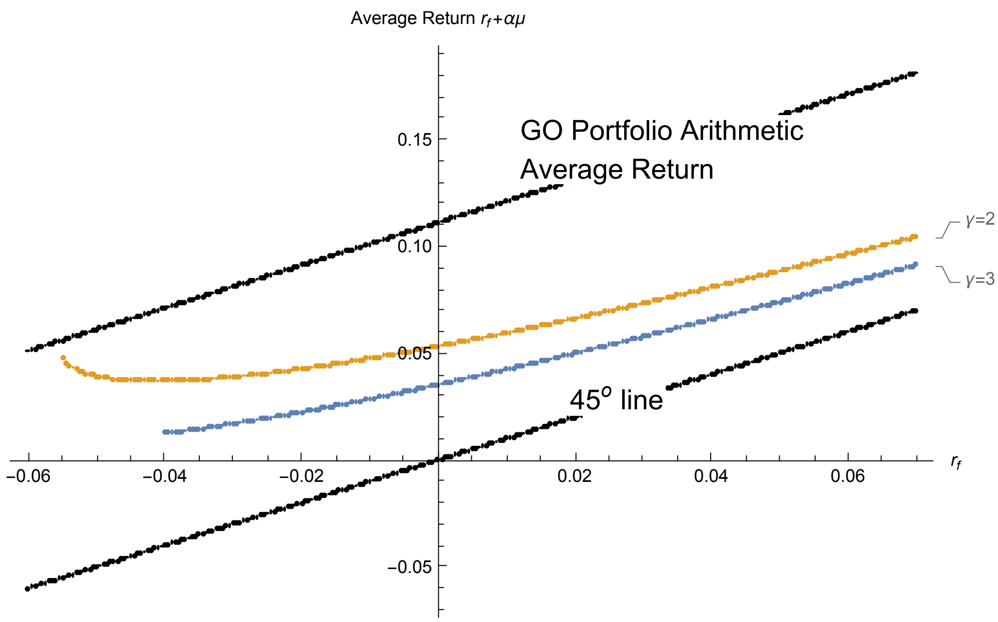

Figure 4: Reaching for Yield. Portfolio Expected Return, the Riskfree Rate, and Risk Aversion.

Existence of a solution Figure 4 also illustrates the fact that two different conditions need to be satisfied for a solution to exist. As we discussed in section 2, the parameters of the problem must allow positive consumption for some feasible portfolio:

$$r_f > -\frac{1}{2}\frac{\mu^2}{\sigma^2}. \tag{39}$$

In addition, the volatility of consumption must not be so large that the value function fails to converge:

$$\alpha < \sqrt{\frac{2\rho}{(1-\gamma)^2\sigma^2}}. \tag{40}$$

Depending on the parameters of the problem, and in particular on risk aversion, either condition (39) or (40) may bind.

Figure 4 shows these two different conditions binding in the two cases plotted. For a relatively aggressive investor with risk aversion $\gamma = 2$, condition (39) binds when the riskfree rate is very low.

Below the binding value of the riskfree rate, even the growth-optimal portfolio does not provide positive consumption. We can see this in Figure 4 where the upper orange curve starts from the growth-optimal optimal portfolio line.[13]

For a more conservative investor with risk aversion $\gamma = 3$, on the other hand, condition (40) binds. We can see this in Figure 4, where the leftmost point of the lower blue curve is to the right of the leftmost point of the upper curve. A conservative investor with a geometric average spending constraint cannot tolerate as low a riskfree interest rate as can a more aggressive investor with the same constraint.

The effect of the risk premium on risktaking In the standard Merton model a larger risk premium unambiguously raises the share of wealth allocated to risky assets. Proposition 2 shows that this is no longer the case in the model with a geometric sustainable spending constraint. In Figure 5 we plot the optimal risky share as a function of the riskfree rate for different values of the risk premium.[14] We see that the curves turn counter-clockwise around a particular point with a negative riskfree interest rate. At this point the risky share is independent of the risk premium. At higher values of the riskfree rate, the risk premium has the usual positive effect on the risky share. But at lower (more negative) values of the riskfree rate the risk premium has a negative effect on the risky share.

To understand this result we note that, as in the arithmetic average model, there are two effects in play when the risk premium changes. An increase in the risk premium both shifts the portfolio constraint up, and makes it steeper. The former effect reduces risktaking and the latter increases it. As the riskfree rate declines, the former reaching for yield effect becomes more powerful and it dominates at sufficiently low levels of the riskfree rate.

[13] Since the figure plots the arithmetic average portfolio return, the line is positive at this point but the geometric average portfolio return that enters the sustainable spending constraint is zero.

[14] All parameters are the same as in Figure 1 except that we consider risk premia of 5%, 6%, and 7%.

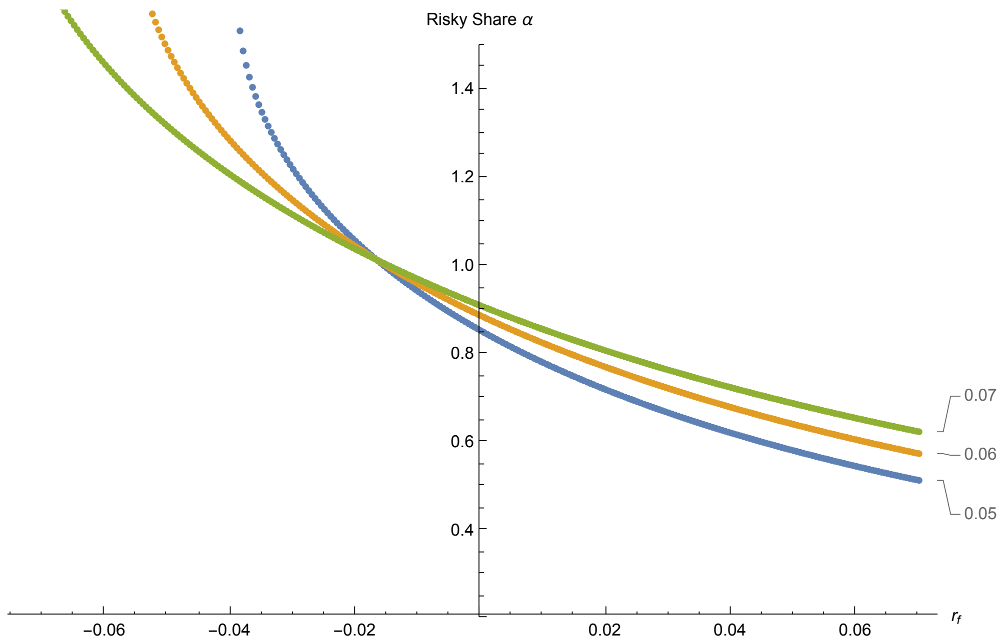

Figure 5: Sensitivity of Optimal Risky Share to Risk Premium

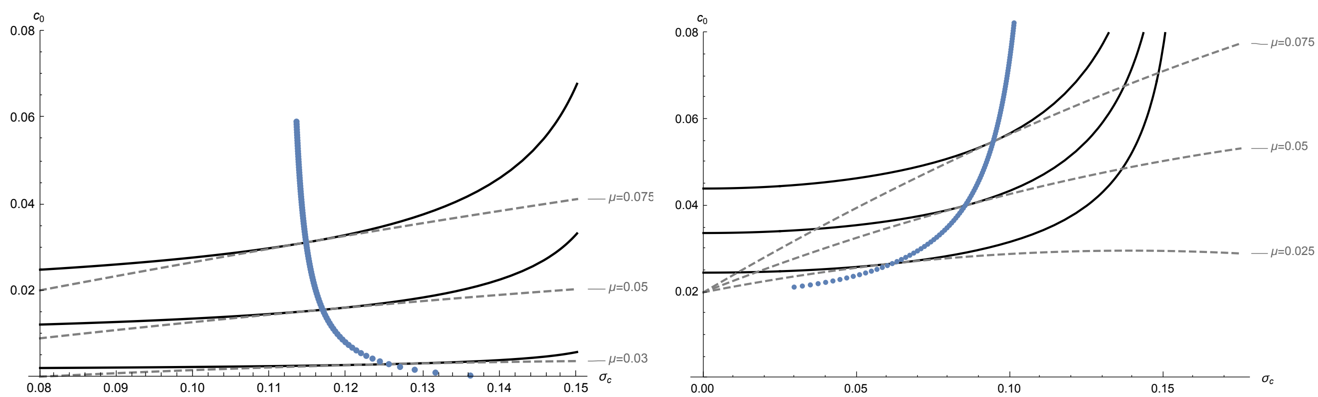

Figure 6: The Effect of Risk Premium on Initial Consumption and Consumption Volatility

Figure 6 has two panels illustrating the two possible situations.[15] Panel A shows a case where r_f is very low at -1%. Here an increase in value leads to a significant steepening of the iso-value curve. This effect dominates the steepening of the portfolio constraint and results in lower consumption volatility for larger values of the risk premium. Panel B shows the case where r_f is 2%. Here an increase in value results in an almost parallel shift of the iso-value curves that is dominated by the steepening of the portfolio constraint.

5 Time-Varying Riskfree Rate

Earlier sections of this paper have compared static equilibria with higher or lower real interest rates that are assumed to be constant over time. It is natural to ask whether similar results apply in a dynamic model where the riskfree interest rate moves over time.

In a dynamic environment where investment opportunities change over time, the demand for a risky asset depends on its intertemporal hedging properties (Merton 1973). As in Campbell and Viceira (2001), we allow the riskfree interest rate to move but assume a constant risk premium. Then, intertemporal hedging is driven by correlation between shocks to the return on the risky asset and shocks to the riskfree interest rate. We capture such correlation by allowing interest rate innovations to load on two independent Brownian motions, one of which also drives the return on the risky asset. To see how this works, assume two independent Brownian motions dZ_{1t} and dZ_{2t}. Then a third Brownian motion $dZ_{3t} = \eta dZ_{1t} + \sqrt{1-\eta^2} dZ_{2t}$ has instantaneous variance dt and instantaneous correlation ηdt with dZ_{1t}.

The full problem of the agent now becomes

$$\max_{\alpha_t} E_0 \int_0^\infty e^{-\rho t} u(c_t) dt \tag{41}$$

[15] Risk aversion, the rate of time preference, and the volatility of the risky asset return are the same as in Figure 1.

$$\text{subject to } c_t = w_t\left(r_{ft} + \alpha_t\mu - \frac{1}{2}\alpha_t^2\sigma^2\right)$$
$$\begin{pmatrix} dw_t \\ dr_{ft} \end{pmatrix} = \begin{pmatrix} \frac{1}{2}w_t\alpha_t^2\sigma^2 \\ \phi(r_{ft}) \end{pmatrix} + \begin{pmatrix} w_t\alpha_t\sigma & 0 \\ \nu r_{ft}\eta & \nu r_{ft}\sqrt{1-\eta^2} \end{pmatrix}\begin{pmatrix} dZ_{1t} \\ dZ_{2t} \end{pmatrix} \tag{42}$$

HJB equation Under these assumptions the HJB equation is

$$\rho v(w_t, r_{ft}) = \max_{\alpha}\left\{u(c_t) + \frac{\partial v}{\partial w}\frac{1}{2}w_t\alpha^2\sigma^2 + \frac{\partial v}{\partial r_f}\phi(r_{ft}) + \frac{1}{2}\frac{\partial^2 v}{\partial w^2}w_t^2\alpha^2\sigma^2 + \frac{1}{2}\frac{\partial^2 v}{\partial r_f^2}\nu^2 r_{ft}^2 + \frac{\partial^2 v}{\partial w\partial r_f}w_t\alpha\nu r_{ft}\eta\right\}. \tag{43}$$

As previously, we conjecture that the value function takes the form

$$v(w_t, r_{ft}) = A(r_{ft})\frac{w_t^{1-\gamma}}{1-\gamma}, \tag{44}$$

which differs from the static model by having a time-varying coefficient $A(r_{ft})$ that reflects changing investment opportunities. If we substitute this guess for the value function into the HJB equation, wealth cancels thus verifying the conjectured form of the value function. We collect terms to obtain

$$\begin{aligned}\rho A(r_{ft})\frac{1}{1-\gamma} = \max_{\alpha}\frac{\left(r_{ft} + \alpha\mu - \frac{1}{2}\alpha^2\sigma^2\right)^{1-\gamma}}{1-\gamma} &+ A(r)\frac{1}{2}(1-\gamma)\alpha^2\sigma^2 + A'(r_{ft})\frac{1}{1-\gamma}\phi(r_{ft}) \\ &+ \frac{1}{2}A''(r_{ft})\frac{1}{1-\gamma}\nu^2 r_{ft}^2 + A'(r_{ft})\alpha\nu r_{ft}\eta.\end{aligned} \tag{45}$$

First-order condition The first order condition for this problem has an additional term compared to the static model:

$$\left(r_{ft} + \alpha\mu - \frac{1}{2}\alpha^2\sigma^2\right)^{-\gamma}(\mu - \alpha\sigma^2) = A(r_{ft})(\gamma-1)\alpha\sigma^2 - A'(r_{ft})r_{ft}\nu\eta. \tag{46}$$

The first-order condition shows that α depends on the correlation between the unexpected risky asset return and the risk free rate η. A lower correlation η decreases the right-hand side of equation (46) since $A'(r_{ft}) < 0$. The left-hand side is decreasing in α, so the optimal risky share α is larger for a smaller correlation η.

The drift of log consumption In the static model the geometric average sustainable spending rule implies that the logs of both wealth and consumption are martingales. In the dynamic model, log wealth remains a martingale but log consumption is not in general. To see this, consider the log change in consumption

$$\begin{aligned} d\log c_t &= d\log w_t \left(r_{ft} + \alpha(r_{ft})\mu - \frac{1}{2}\alpha(r_{ft})^2\sigma^2 \right) \\ &= \underbrace{d\log w_t}_{\text{mean zero}} + d\log\left(r_{ft} + \alpha(r_{ft})\mu - \frac{1}{2}\alpha(r_{ft})^2\sigma^2 \right). \end{aligned} \tag{47}$$

In general the second term has nonzero drift and we cannot characterize it analytically. Therefore, we rely on numerical methods to describe this and other properties of the solution. The online appendix provides more details on the numerical implementation.

Analysis of the numerical solution In Figure 7, we present a numerical solution for the dynamic model and illustrate its sensitivity to different parameter values. We set

$$\phi(r_{ft}) = \frac{1}{2}\nu^2 r_{ft} \tag{48}$$

in order to make $\log(r_{ft})$ a random walk. This process is natural for our application since it implies that log consumption for an agent that takes zero risk ($\alpha = 0$) follows a random walk similarly to the static case. Log consumption for the agent who invests optimally is not quite a random walk, but predictable changes in consumption are small. More generally, a persistent process for

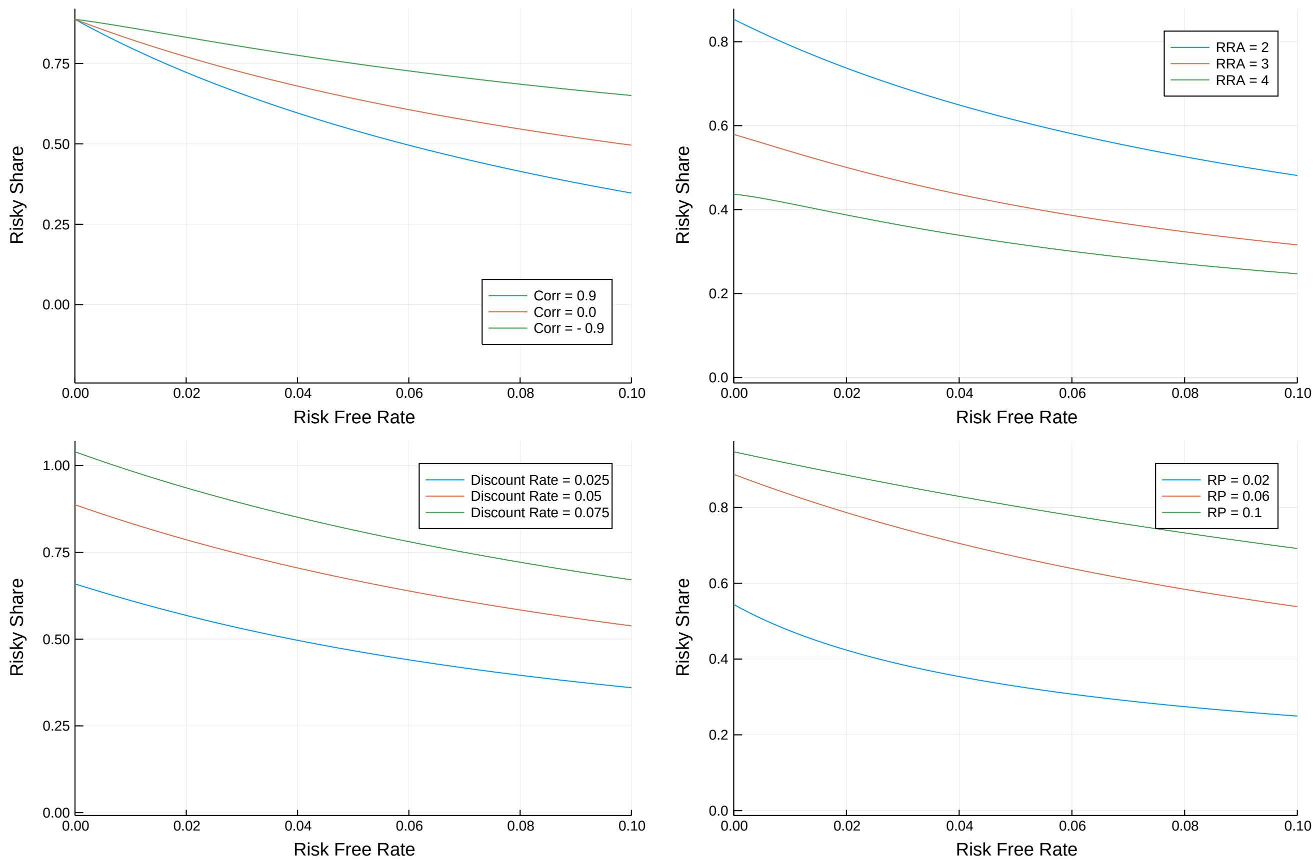

Figure 7: Sensitivity of Optimal Risky Share to Parameters in Dynamic Model

the riskfree rate is required to avoid transient movements in consumption and asset allocation that would be generated by a sustainable spending constraint interacting with transitory shocks to the expected return.

All four panels of Figure 7 have the riskfree interest rate on the horizontal axis, and the risky share on the vertical axis. All the panels show decreasing convex curves, consistent with the static analysis of reaching for yield. The panels differ in the parameters assumed. In every panel the base case, shown as a red curve, has a zero correlation between the riskfree rate and the risky asset return, risk aversion of 4, a time discount rate of 5%, a risk premium of 6%, and a risky standard deviation of 18% implying a Sharpe ratio of 1/3.

The top left panel varies the correlation to 0.9 or -0.9. A negative correlation increases the demand for the risky asset while a positive correlation reduces it, consistent with the logic of Merton (1973, 1993) and Campbell and Viceira (2001). The top right panel varies risk aversion up to 5 or down to 3, with intuitive effects on risktaking. The bottom left panel varies the rate of time preference up to 7.5% or down to 2.5%, and as in the static model a more impatient investor takes more risk. The bottom right panel varies the risk premium up to 10% or down to 2%. Since in this model the riskfree interest rate is always positive, the effect of the risk premium on risktaking is always positive, but the magnitude of the effect is smaller when the interest rate reaches very low levels. These findings confirm the main results of our static analysis for a plausible dynamic model.

6 Conclusion

In this paper we have shown that a constraint on an investor's ability to save or dissave can break the standard result that risktaking depends only on risk aversion, risk, and the risk premium available in financial markets. An investor with a sustainable spending constraint reaches for yield, taking more risk as the riskfree interest rate declines, even if all the standard determinants of risktaking are constant. Furthermore, the tendency to reach for yield is stronger when the real interest rate is low than when it is high. This may be one reason why reaching for yield has been so widely discussed in the low-interest-rate environment of the early 21st Century.

Reaching for yield also changes the investor's response to a change in the risk premium. An increase in the risk premium stimulates risktaking through the conventional channel, a substitution effect towards risky investing. However it also weakens reaching for yield, an offsetting income effect that reduces risktaking. The offsetting effect becomes stronger when the riskfree interest rate is low, and can even dominate at sufficiently low levels of the interest rate.

Most of our analysis is static, comparing equilibria with permanently different levels of the riskfree interest rate. However we show that our results carry over to the numerical solution of a dynamic model with a highly persistent riskfree rate whose log follows a random walk.

Although our results depend on the functional forms we have assumed—both the power form for utility and the arithmetic or geometric forms for the sustainable spending constraints—we believe the insights of this paper are more general. Most obviously, it should be straightforward to extend the analysis to Epstein-Zin preferences, breaking the link between risk aversion and the elasticity of intertemporal substitution implied by power utility.

It should also be possible to model a sustainable spending inequality that allows saving (consuming less than the expected return on wealth) but not dissaving. Such an inequality would bind only at low levels of interest rates that would otherwise encourage dissaving, providing an additional reason why reaching for yield becomes more important when the riskfree interest rate is low. It could be augmented by a fixed cost of adjusting the spending rate (the ratio of spending to wealth), consistent with the fact that endowments typically adjust their spending rates discretely rather than continuously.

Finally, we believe that more flexible constraints, for example constraints that penalize but do not prohibit saving or dissaving, will also imply some degree of reaching for yield behavior. In this sense the classic result that the riskfree interest rate has no effect on risktaking may be regarded as a knife-edge case that often fails to describe reality.

References

Acharya, Viral, and Hassan Naqvi, 2019, "On Reaching for Yield and the Coexistence of Bubbles and Negative Bubbles", *Journal of Financial Intermediation* 38, 1–10.

Andonov, Aleksandar, Rob M.M.J. Bauer, and K.J. Martijn Cremers, 2017, "Pension Fund Asset Allocation and Liability Discount Rates", *Review of Financial Studies* 30, 2555–2595.

Becker, Bo and Victoria Ivashina, 2015, "Reaching for Yield in the Bond Market", *Journal of Finance* 70, 1863–1901.

Borio, Claudio and Haibin Zhu, 2012, "Capital Regulation, Risk-Taking, and Monetary Policy: A Missing Link in the Transmission Mechanism?", *Journal of Financial Stability* 8, 236–251.

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, 931–62.

Campbell, John Y., 2018, *Financial Decisions and Markets: A Course in Asset Pricing*, Princeton University Press, Princeton, NJ.

Campbell, John Y. and Roman Sigalov, 2020, "Appendix to Portfolio Choice with Sustainable Spending: A Model of Reaching for Yield", https://scholar.harvard.edu/campbell/publications.

Campbell, John Y. and Luis M. Viceira, 2001, "Who Should Buy Long-Term Bonds?", *American Economic Review* 91, 99–127.

Chodorow-Reich, Gabriel, 2014, "Effects of Unconventional Monetary Policy on Financial Institutions", *Brookings Papers on Economic Activity* (Spring), 155–204.

Choi, Jaewon, and Mathias Kronlund, 2017, "Reaching for Yield in Corporate Bond Mutual Funds", *Review of Financial Studies* 31, 1930–1965.

Di Maggio, Marco, and Marcin Kacperczyk, 2017, "The Unintended Consequences of the Zero Lower Bound Policy", *Journal of Financial Economics* 123, 59–80.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2018, "A Model of Monetary Policy and Risk Premia", *Journal of Finance* 73, 317–373.

Dybvig, Philip H., 1995, "Duesenberry's Ratcheting of Consumption: Optimal Dynamic Consumption and Investment Given Intolerance for Any Decline in Standard of Living", *Review of Economic Studies* 62, 287–313.

Dybvig, Philip H. and Zhenjiang Qin, 2019, “How to Squander Your Endowment: Pitfalls and Remedies”, unpublished paper, Washington University in St. Louis and University of Macao.

Gilbert, Thomas and Christopher Hrdlicka, 2015, “Why Are University Endowments Large and Risky?”, *Review of Financial Studies* 28, 2643–2686.

Hanson, Samuel G. and Jeremy C. Stein, 2015, “Monetary Policy and Long-Term Real Rates”, *Journal of Financial Economics* 115, 429–448.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina, 2014, “Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say About the Effects of Monetary Policy on Credit Risk-Taking?”, *Econometrica* 82, 463–505.

Lian, Chen, Yueran Ma, and Carmen Wang, 2019, “Low Interest Rates and Risk-Taking: Evidence from Individual Investment Decisions”, *Review of Financial Studies* 32, 2107–2148.

Lu, Lina, Matthew Pritsker, Andrei Zlate, Ken Anadu, and James Bohn, 2019, “Reach for Yield by U.S. Public Pension Funds”, unpublished paper, Federal Reserve Bank of Boston and Board of Governors of the Federal Reserve System.

Maddaloni, Angela, and José-Luis Peydró, 2011, “Bank Risk-Taking, Securitization, Supervision, and Low Interest Rates: Evidence from the Euro-Area and the US Lending Standards”, *Review of Financial Studies* 24, 2121–2165.

Martin, Ian W.R., 2012, “On the Valuation of Long-Dated Assets”, *Journal of Political Economy* 120, 346–358.

Merton, Robert C., 1969, “Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case”, *Review of Economics and Statistics* 51, 247–257.

Merton, Robert C., 1971, “Optimum Consumption and Portfolio Rules in a Continuous-Time Model”, *Journal of Economic Theory* 3, 373–413.

Merton, Robert C., 1973, “An Intertemporal Capital Asset Pricing Model”, *Econometrica* 41, 867–887.

Merton, Robert C., 1993, “Optimal Investment Strategies for University Endowment Funds”, in Charles T. Clotfelter and Michael Rothschild, eds. *Studies of Supply and Demand in Higher Education*, University of Chicago Press, 211–242.

Rajan, Raghuram, 2006, “Has Finance Made the World Riskier?”, *European Financial Management* 12, 499–533.

Rajan, Raghuram, 2013, "A Step in the Dark: Unconventional Monetary Policy After the Crisis", Andrew Crockett Memorial Lecture, Bank for International Settlements. Available online at https://www.bis.org/events/agm2013/sp130623.htm.

Sitkoff, Robert H. and Jesse Dukeminier, 2017, *Wills, Trusts, and Estates*, 10th edition, Wolters Kluwer, New York, NY.

Stein, Jeremy, 2013, "Overheating in Credit Markets: Origins, Measurement, and Policy Responses", Research Symposium, Federal Reserve Bank of St. Louis. Available online at https://www.federalreserve.gov/newsevents/speech/stein20130207a.htm.