International Review of Financial Analysis 88 (2023) 102696

Contents lists available at ScienceDirect

International Review of Financial Analysis

journal homepage: www.elsevier.com/locate/irfa

Expected long-term rates of return when short-term returns are serially correlated☆

Knut Anton Mork *, Haakon Andreas Trønnes

Norwegian University of Science and Technology (NTNU), Department of Economics, Norway

ARTICLE INFO

JEL classification:
C61
G11
G15

Keywords:
Serially correlated rates of return
Purchasing power parity
Mean reversion
Withdrawal rules

ABSTRACT

When short-term returns are serially uncorrelated, expected long-term and short-term returns are equal. However, we show that negative serial correlation among the short-term returns make the expected long-term returns lower than the short-term ones. Such serial correlation is likely to arise, for example, for an investor whose portfolio is invested abroad in assets denominated in foreign currencies, but who wants to make withdrawals in proportion to the fund's value in the domestic currency, and the exchange rate obeys long-term purchasing-power parity. Small-country sovereign wealth funds are leading examples of such investors. For the Norwegian GPFG, the expected annualized long-term rate of return in Norwegian kroner may be 0.7 to 1.8 percentage points lower than the expected short-term return. Negative contemporaneous correlation between global returns and changes in the real exchange rate may dampen and even reverse this result. Empirical evidence suggests that this may be the case for investors in USD-denominated assets domiciled in the United Kingdom, but not for investors domiciled in Norway or Germany. Empirically, we furthermore find that long-term annualized returns may fall short of short-term returns even when evaluated in real USD. Although negative serial correlation also shrinks the long-term variance, funds such as the GPFG should calibrate withdrawals to the expected long-term returns rather than the short-term ones.

1. Introduction

The owners of many endowment funds and sovereign wealth funds have implemented rules specifying that regular withdrawals be tied to the fund's expected real return. This return is typically estimated as an average of past annual returns. However, if the rates of return display negative serial correlation, such calculations typically underestimate the average return that can be expected over long horizons. The bias can be substantial and lead to premature fund depletion. A prominent example is the Norwegian Government Pension Fund Global (GPFG), the world's largest sovereign wealth fund with a current value of about USD 1.3 trillion, which inspired this paper. Withdrawals from this fund are regulated by the kind of rule just mentioned with the intention of forever preserving the fund's real value in expectation. However, we estimate that the typical method used to calculate the annual real return underestimates the long-term return by 0.7 to 1.8 percentage point per year. Considering that the rule for spending of the proceeds of this fund specifies annual withdrawals up to amounts corresponding to the expected one-year rate of return, these results imply a rather strong warning that the fund may be depleted in finite time.

This result is related to the findings in the literature of skewness in long-term returns. Arditti and Levy (1975) show that compounding of risky returns induces rightward skewness in multi-period returns even if single-period returns are symmetrical. Bessembinder (2018) finds strong empirical support for this result. The skewness rises with the time time horizon and the volatility of the return. Consider, as an example, an asset that each period has an equal chance of rising or falling 50% in value. If it is held for one period only, both mean and median return are zero. If it is held for two periods, the mean remains zero, but the median drops to an annualized – 13.3%. This skewness of the return distribution is not an artifact of this particular example, but a general implication of compounding.

Serial correlation of returns adds a layer of complication, which is the subject of this paper. Suppose that, in the second period of holding,

☆ We have benefitted from useful comments by Fabian A. Harang, Knut K. Aase, Snorre Lindset, and seminar participants at the Norwegian Ministry of Finance. The responsibility for errors and omissions is solely our own. We have benefitted from financial assistance from Finansmarkedsfondet, the financial markets research fund under The Research Council of Norway, Project no. 294398.

* Corresponding author.
E-mail addresses: knut.anton.mork@ntnu.no (K.A. Mork), haakon.a.tronnes@ntnu.no (H.A. Trønnes).

https://doi.org/10.1016/j.irfa.2023.102696
Received 10 January 2023; Received in revised form 11 May 2023; Accepted 22 May 2023
Available online 26 May 2023

the probability of a 50% gain is 40% if the first-period return was 50%, but 60% if the first-period return was – 50%. The probability of a 50% loss is similarly 60% if the first-period return was 50%, and 40% if it was – 50%. This change does not affect the unconditional probability distribution for the second-period return, whose mean remains zero. The annualized median two-period return remains −13.3% as well. However, the *mean* two-period return drops to −2.5%. The negative serial correlation tightens the distribution from both ends. However, because the high end is further off from the median than the low end, the tightening on the high end has a greater effect on the expectation.

This shrinkage of the annualized multiperiod return does not mean that the second-period return is expected to be lower than the first-year one, conditional on the information at the time the security is acquired. Again, it is also not an artifact of the example chosen, nor does it come from any abuse of Itô's lemma, although we do make use of that well-known result below. It happens simply because the compounding of returns is not a linear operation, but an inherently convex one.

Positive serial correlation has the opposite effect, but negative serial correlation seems more likely to be observed in practice. Fama and French (1988) claimed that stock prices show evidence of mean reversion when studied in long enough samples. The ensuing debate about their results does not seem to have been settled in the literature, cf. Poterba and Summers (1988), Mukherji (2011), and Pástor and Stambaugh (2012). However, the recent essay on long-term investment by Cochrane (2022) devotes considerable attention to the possibility that stock price drops driven by increases in investors' discount rates ("discount rate betas") may signal higher returns ahead. Cochrane emphasizes the similarity between such movements in stock returns and those of long-term bonds, which indeed tend to be negatively serially correlated, cf. e.g., Campbell and Viceira (2002).

Another, more clear-cut example concerns investors, like the owners of the Norwegian GPFG, that hold foreign-currency denominated assets but consume goods whose prices are denominated in the domestic currency. For such investors, the period real rate of return is a sum of two components, one of which being the real return in the foreign market and the other the period rate of change of the real exchange rate. If the exchange rate obeys long-term purchasing parity, this second component of the real rate of return will be negatively serially correlated.

To our knowledge, this financial effect of purchasing power parity has not been noted in the existing literature. Whereas purchasing power parity may not hold universally, we believe it tends to hold for most developed-country currencies. Chortareas and Kapetanios (2009) provide strong empirical evidence of mean reversion in real exchange rates in the OECD. A further review is provided by Rabe & Waddle (2020). Hebisha (2023) finds evidence of mean reversion for most of the currencies that were converted to the euro in 1999. Our attention to this phenomenon came from our research on the Norwegian GPFG, whose entire portfolio is held in foreign assets denominated in foreign currencies, but whose returns are used to help fund the government's domestic spending, which naturally is denominated in the domestic currency.

Currency exchange adds an element of risk to the small-country investor in the global market. The negative serial correlation actually limits this risk; and the convexity introduced by compounding makes sure that the limitation happens mainly on the upside. However, that is enough to bias the expected fund value downwards.

A further limitation of risk arises if the investor country's real exchange rate is negatively correlated with the movements in global financial markets. The presence of such contemporaneous correlation dampens the bias just mentioned and may even reverse its sign if the contemporaneous correlation is strong enough.

The following pages explain how that happens and presents some evidence on the empirical importance of the relevant mechanisms. The next section presents a simple example in discrete time. In Section 3, we develop a rigorous model in continuous time for the annualized expected long-term returns for a portfolio whose dynamic behavior can be described as the sum of a submartingale and a mean-reversion process. Section 4 adds the possibility of a negative contemporaneous correlation between the two elements and shows how this correlation may dampen or reverse the results in Section 3. Section 5 presents some empirical evidence on the expected multi-period real returns in local currencies of dollar-based portfolios owned by investors domiciled in Norway, Germany, and the United Kingdom, respectively. Section 6 discusses the results, and Section 7 concludes.

2. An example in discrete time

Let r_t denote the annual rate of return in from year $t-1$ to year t. Assume the return process is stationary with mean $\bar{r}$ and serial covariances $cov(r_t, r_{t-j})$, denoted s_j. Then, as shown in Appendix A, the annualized expected return of holding this asset from year 0 to year $n \geq 2$ can be written as

$$\bar{r}_{0n} = (1+\bar{r})\left\{1+(1+\bar{r})^{-2}\left[\sum_{j=1}^{n-1}(n-j)s_j + \text{h.o.m.}\right]\right\}^{1/n} - 1, \tag{1}$$

where h.o.m. stands for higher order moments. Clearly, if all the serial correlation terms are weakly negative, some of them strictly negative, and the higher-order terms small enough to be ignored, $\bar{r}_{0n} < \bar{r}$.

The source of this effect lies in the compounding of returns. Whereas the financial returns are compounded geometrically, the expected value of the rate of return is computed arithmetically. If one year's rate of return comes out above expectation it will make the fund larger. Because of the negative serial correlation, this movement will, with some probability, be reversed the following year. The same mechanism works on the downside. However, the reversal from the upside will be larger because it starts from a higher base. Thus, on average, the lingering effects of a less-than-expected rate of return will dominate those of the larger-than-expected ones.

3. Formal analysis in continuous time

We consider an asset or a portfolio of assets whose rate of return contains an element of negative serial correlation. In this setting, the rate of return can conveniently be specified as a sum of two components, a geometric Brownian motion with drift μ and standard deviation σ_1, and a mean-reverting Ornstein-Uhlenbeck process $dx(t)$, so that the instantaneous return r on the asset (or fund) value A is given by the sum

$$r(t) := \frac{dA(t)}{A(t)} = \mu dt + \sigma_1 dB_1(t) + dx(t), \tag{2}$$

where the last term is the Ornstein-Uhlenbeck process

$$dx(t) = \alpha(\kappa - x(t))dt + \sigma_2 dB_2(t). \tag{3}$$

Here, B_2 is a Brownian motion, stochastically independent of B_1, and $\alpha > 0$ and κ are constant parameters. This process, originally proposed by the physicists Uhlenbeck and Ornstein (1930) and explored by, for example, Vasicek (1977) and Maller, Müller, G., & Szimayer (2009),[1] is the continuous-time analogy of an AR(1) process in discrete time. We emphasize that we think of the two components of the rate of return as representing two properties of the same asset or portfolio, not as different assets or portfolio components among which the investor can choose. From (3) it should be clear that this component of the return has an unconditional zero mean, so that $\mathbb{E}r(t) = \mu$.

The differential equation Eq. (3) is known to have the following solution:

[1] A compact presentation can be found in Wikipedia, https://en.wikipedia.org/wiki/Ornstein%E2%80%93Uhlenbeck_process, accessed on June 25, 2022.

$$x(t) = x_0 e^{-\alpha t} + \kappa(1 - e^{-\alpha t}) + \sigma_2 \int_0^t e^{-\alpha(t-s)} dB_2(s). \tag{4}$$

As shown in Appendix B, the unconditional (i.e. t, $s \to \infty$) serial correlation coefficients for the rates of return in (2) can be written as

$$corr(r_t, r_s) = -\frac{\alpha}{2} e^{-\alpha|t-s|} \left(\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \right) dt, t \neq s. \tag{5}$$

Clearly, $corr(r_t, r_s) < 0$. Its absolute value is lower the longer the lag length $|t-s|$ and approaches zero asymptotically. For a given lag length, it increases with the mean-reversion parameter α and also with the relative importance of the mean-reverting component of the rate of return, as measured by the variance ratio $\sigma_2^2/(\sigma_1^2 + \sigma_2^2)$.

Substituting from (3) into (2), we can write the instantaneous return as:

$$r(t) = \frac{dA(t)}{A(t)} = [\mu + \alpha(\kappa - x(t))]dt + \sigma_1 dB_1(t) + \sigma_2 dB_2(t). \tag{6}$$

Proposition. *For an asset or portfolio with a starting value of $A_0$and instantaneous rates of return defined by (6), the annualized expected return of holding the asset or portfolio until time $t > 0$ is (a) lower than the expected short-term return, (b), a decreasing function of the length t of the holding period, (c) asymptotically reaching a lower limit equal to the annualized expected rate of return for the limiting case where x(t) is a Brownian motion, i.e. $\alpha \to \infty$.*

Proof. Application of Itô's lemma to (6) yields the law of motion in logs as

$$dlnA(t) = \left[\mu + \alpha(\kappa - x(t)) - \frac{1}{2}(\sigma_1^2 + \sigma_2^2) \right] dt + \sigma_1 dB_1(t) + \sigma_2 dB_2(t)$$

$$= \left[\mu - \frac{1}{2}(\sigma_1^2 + \sigma_2^2) \right] dt + \sigma_1 dB_1(t) + dx(t). \tag{7}$$

Integration forward then gives the future log value as

$$ln\left(\frac{A(t)}{A_0}\right) = \left[\mu - \frac{1}{2}(\sigma_1^2 + \sigma_2^2) \right] t + \sigma_1 \int_0^t dB_1(s) + \int_0^t dx(s)$$

$$= \left[\mu - \frac{1}{2}(\sigma_1^2 + \sigma_2^2) \right] t + \sigma_1 \int_0^t dB_1(s) + x(t) - x_0$$

$$= \left[\mu - \frac{1}{2}(\sigma_1^2 + \sigma_2^2) \right] t + \sigma_1 \int_0^t dB_1(s) + (\kappa - x_0)(1 - e^{-\alpha t}) + \sigma_2 \int_0^t e^{-\alpha(t-s)} dB_2(s), \tag{8}$$

where the last equality follows from (4). From (4), we also find the unconditional mean

$$\mathbb{E}[x(t) - x_0] = \mathbb{E}[(\kappa - x_0)(1 - e^{-\alpha t})] = 0,$$

so that

$$\mathbb{E}ln\left(\frac{A(t)}{A_0}\right) = \left[\mu - \frac{1}{2}(\sigma_1^2 + \sigma_2^2) \right] t.$$

Furthermore, from (8):

$$\mathbb{V}ln\left(\frac{A(t)}{A_0}\right) = t\sigma_1^2 + \sigma_2^2 \frac{1 - e^{-2\alpha t}}{2\alpha}. \tag{9}$$

From the formula for the expectation of a lognormal variable, we then find

$$ln\mathbb{E}\left(\frac{A(t)}{A_0}\right) = \mu t - \frac{1}{2}\sigma_2^2 \left(t - \frac{1 - e^{-2\alpha t}}{2\alpha} \right). \tag{10}$$

Thus, the annualized expected return after a holding period of t years becomes

$$\bar{r}_{0t} = \left(\frac{1}{t}\right) ln\mathbb{E}\left(\frac{A(t)}{A_0}\right) = \mu - \frac{1}{2}\sigma_2^2 \left(1 - \frac{1 - e^{-2\alpha t}}{2\alpha t} \right). \tag{11}$$

Note that the function $f : \mathbb{R}^+ \to (0, 1)$, defined by

$$f(z) = \frac{1 - e^{-z}}{z} \tag{12}$$

is convex and monotonically decreasing between the limiting values

$$\lim_{z \to 0} f(z) = 1$$

and

$$\lim_{z \to \infty} f(z) = 0,$$

so that

$$0 < f(z) < 1.$$

Thus, for $\alpha t > 0$,

$$\bar{r}_{0t} = \mu - \frac{1}{2}\sigma_2^2 [1 - f(2\alpha t)] \langle \mu.$$

$\bar{r}_{0t}$ is decreasing in t and α, and approaches a lower limit of $\mu - \sigma_2^2/2$ for $t \to \infty$, the same limit as the one approached for all horizons if $\alpha \to \infty$. **Q.E.D.**

Consider now the case of an investor located in a small country with an independent currency, but whose entire investment is made abroad. Measured in the domestic currency, the real rate of return is then the sum of the real return in foreign currency and the rate of change of the real exchange rate. If the exchange rate satisfies purchasing power parity as a long-term trend, the real exchange rate can reasonably be modeled as a mean-reverting process like the one in (3). Supposing that the first two components on the right of (2) represent the real return in the global market and that the two stochastic processes are independent of each other, we have then proved

Corollary 1. *For an investor located in a country whose real exchange rate behaves like in (3) and whose investment is made abroad at real returns in foreign currency as specified by the first two terms on the right of (2), the annualized expected real return in the domestic currency of holding the asset or portfolio over time is lower than the corresponding real rate in the global market and a decreasing function of the length t of the holding period.*

A number of investment funds, such as the Norwegian GPFG, are subject to rules that permit annual withdrawals in amounts corresponding to the fund's expected real return. In our model, this behavior can be specified by setting $\mu = 0$. Then, if the entire fund is invested abroad, but the withdrawals specified as a percentage of the fund's value in the domestic currency, Proposition 1 means that the fund will eventually be depleted in expectation. We formulate this result as

Corollary 2. *Consider a fund owned by an agent in a country whose real exchange rate behaves like (3), whose investment is made abroad at real returns in foreign currency as given by the first two terms on the right of (2), and from which regular withdrawals are made in amounts equal to μ times the fund's value in the local currency. Then, the fund's expected value in domestic currency is a decreasing function of the time horizon t, asymptotically approaching zero as $t \to \infty$.*

Proof. Setting $\mu = 0$ in (10) gives

$$ln\mathbb{E}\left(\frac{A(t)}{A_0}\right)\Big\langle 0,$$

so that

$$\mathbb{E}A(t) < A_0.$$

Clearly, the right-hand side of (10) then grows without limit in absolute value as $t \rightarrow \infty$, so that $\mathbb{E}A(t) \rightarrow 0$. **Q.E.D.**

Fig. 1 illustrates the results in the proposition by depicting the annualized expected future returns for horizons up to 50 years and varying values of the mean-reversion parameter α. Whereas the graph for $\alpha = 0$ (no serial correlation) is a horizontal line at the short-term expected rate of return of μ, the corresponding graphs for $\alpha > 0$ decline with the time horizon and asymptotically approach the lower bound of $\mu - \sigma_2^2/2$, which would represent the annual expected rate of return at all horizons for $\alpha \rightarrow \infty$. The stronger the mean reversion, the faster does the annual expected return drop with the time horizon.

Although the distribution of the annualized expected return becomes increasingly skewed to the right as the holding period increases, as noted above, low-probability prospects of very high future returns nevertheless keep the expected long-term return anchored at the expected short-term return as long as the rates of return are serially uncorrelated. In our analysis, this is the limiting case of $\alpha = 0$, which yields $\bar{r}_{0t} = \mu$.

What changes this result in the presence of negative serial correlation is the fact that mean reversion of securities prices reduces the long-term variance of returns. That effect reduces the likelihood of very high future returns whereas, on the downside, gross returns of zero remains a lower bound. The net result is to lower the expected value and more so as the time horizon rises.

4. Correlation between the return components

So far, our analysis implicitly assumes that the two return components are uncorrelated. For an investor in a small country, it might be argued that a negative contemporaneous correlation should be expected. For example, a shift in sentiment toward "risk off" might make investors shy away from global stocks as well as small-country currencies, so that small-country currencies should tend to depreciate when global equity markets fall. Hossfeld and MacDonald (2015) present evidence of such correlation for the Australian, New Zealand, and Canadian dollars and the Swedish krona, to some extent also the Norwegian krone, especially so during times of financial distress.[2]

To allow for such correlation, redefine dB_2 as

$$dB_2(t) = \rho dB_1(t) + \sqrt{1-\rho^2}dB_2^{'}(t), \tag{13}$$

so that ρ can be interpreted as the coefficient of contemporaneous correlation between dB_1 and dB_2. If this correlation is negative, the exchange-rate mechanism could act like a natural hedge for the small-country investor. Although this correlation leaves the unconditional expectation of the exchange-rate change at zero, it reduces the short-term variance from $\sigma_1^2 + \sigma_2^2$ to $\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2$. This variance could conceivably be lower than σ_1^2, which represents the risk facing a large-country investor with the same portfolio. In that case, the exchange-rate mechanism would serve as an effective natural hedge for the small-time investor. It obviously requires

$$\rho \leq -(1/2)\sigma_2/\sigma_1. \tag{14}$$

However, a negative contemporaneous correlation would change the results in Section 3 even if (14) is not satisfied. With non-zero correlation, Eq. (7) would change into

$$dlnA(t) = \left[\mu + \alpha(\kappa - x(t)) - \frac{1}{2}\left(\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2\right)\right]dt$$

$$+(\sigma_1 + \rho\sigma_2)dB_1(t) + \sigma_2\sqrt{1-\rho^2}dB_2^{'}(t)$$

$$= \left[\mu - \frac{1}{2}\left(\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2\right)\right]dt + \sigma_1 dB_1(t) + dx(t), \tag{7'}$$

and (8) into

$$ln\left(\frac{A(t)}{A_0}\right) = \left[\mu - \frac{1}{2}\left(\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2\right)\right]t + (\kappa - x_0)\left(1 - e^{-\alpha t}\right)$$

$$+\sigma_1\int_0^t dB_1(s) + \rho\sigma_2\int_0^t e^{-\alpha(t-s)}dB_1(s) + \sigma_2\sqrt{1-\rho^2}\int_0^t e^{-\alpha(t-s)}dB_2^{'}(s). \tag{8'}$$

This formula lets us derive the variance of the t-period log change for this case as

$$\mathbb{V}ln\left(\frac{A(t)}{A_0}\right) = t\left[\sigma_1^2 + \sigma_2^2 f(2\alpha t) + 2\rho\sigma_1\sigma_2 f(\alpha t)\right], \tag{9'}$$

where f again is defined as in (12). Then, Eq. (11) becomes

$$\bar{r}_{0t} = \mu - \frac{1}{2}\sigma_2^2[1 - f(2\alpha t)] - \rho\sigma_1\sigma_2[1 - f(\alpha t)]. \tag{11'}$$

Because $f(z) < 1$, the last term in obviously positive if $\rho < 0$. Thus, a negative correlation dampens the results in Section 3 and may indeed reverse them. To see the conditions for such reversal to happen, define the function $g : \mathbb{R}^+ \rightarrow (1,2)$ as

$$g(z) = \frac{1 - f(z)}{1 - f(z/2)}. \tag{15}$$

Like f, g is a convex and decreasing function, with end points are given as

$$\lim_{z\rightarrow 0} g(z) = 2, \lim_{z\rightarrow\infty} g(z) = 1,$$

so that

$$1 < g(z) < 2.$$

For $\bar{r}_{0t} > \mu$, we must have

$$\rho < -\frac{1}{2}(\sigma_2/\sigma_1)g(2\alpha t). \tag{14'}$$

This condition is a little stricter than the one in (14). For a given ρ, it is more likely to be satisfied for long than for short horizons.

Asymptotically, then, the criterion for a contemporaneous negative correlation between the global asset return and the exchange-rate change to reverse the results in Corollaries 1 and 2 is the same as for the currency change to act as an effective hedge in the one-period case. However, because g is a decreasing function, the criterion is somewhat stricter for short than for long horizons.

5. Empirical evidence

As a check for how well our theoretical findings fit the data, we have studied the historical real performance of three different USD-valued portfolios for investors in Norway, Sweden, Germany, and the United Kingdom between January of 1993 and September of 2020. We decided to start our sample at the end of the ERM crisis of the early 1990s, after

[2] Cho, Choi, Kim, and Kim (2016) document flight to quality from emerging to advanced economies during times of stock-market decline. However, because the evidence of purchasing power parity is weaker for emerging-market currencies, we focus our analysis on small, open advanced economies. Ning (2010) analyses the correlation between exchange rates and stock returns in advanced economies; but because she focuses on exchange-rate changes and stock returns in the same country, her results are less relevant for our analysis.

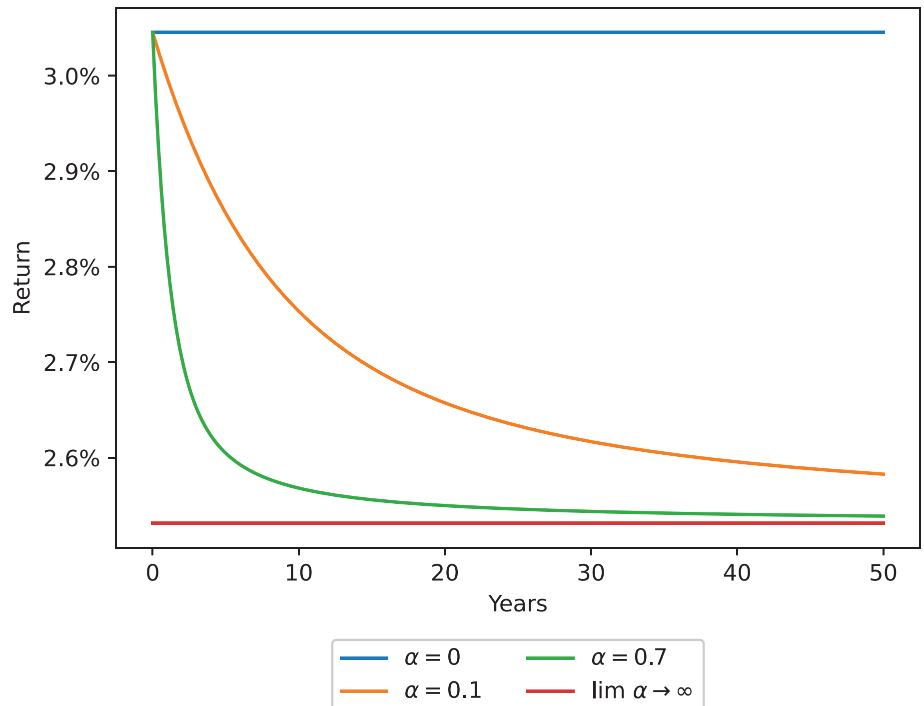

Fig. 1. Annualized expected rate of return by investment horizon.
$\sigma_1 = 17\ \%$, $\sigma_2 = 10\ \%$, $\mu = 3\%$.

which the relevant currencies floated relatively freely.[3] For German investors, we used the D-mark through 1998 and the euro thereafter, splicing the two series at the official rate of conversion. The portfolios were the S&P 500, the FTSE Global Equity, and a 70–30 portfolio of the FTSE Global Equity and the Bloomberg Global Bond index. We included the latter even though bond portfolios may have an inherent tendency toward mean reversion, because this combination approximates the index underlying the government mandate for the Norwegian GPFG. The rates of return include dividends and coupon payments. All indices and exchange rates were deflated with the CPIs of the respective countries.[4]

We started by running Dickey-Fuller tests of unit roots for the respective indices and real exchange rates as a preliminary check of mean reversion. We based our tests on regressions of the form

$$\Delta y_t = c + \delta y_{t-1} + e_t.$$

Our hypothesis is that the three real USD-valued portfolio indices should not have unit roots, implying $\delta = 0$, but have positive constants reflecting mean returns, thus effectively being submartingales. In contrast, the four real exchange rates should be mean reverting, implying $\delta < 0$ and zero constants, as indications of long-run purchasing-power parity. The results are presented in Table 1, with the top panel showing results using monthly data and the lower panel the corresponding results from annual data.

The monthly results find significantly positive constant terms for all three real USD-valued portfolios. They are also consistent with the presence of unit roots, suggesting no mean reversion. The corresponding results on annual data show the same pattern, albeit with somewhat weaker significance for the constant terms. We conclude that these tests show no evidence against the real USD-valued indices behaving like submartingales.

On monthly data, none of the four real exchange rates show significant evidence of mean reversion when constants are included in the regressions. However, the significance of the constant term for the real USD/SEK exchange rate on the 10% level provides some evidence against mean reversion for this currency. Because these constant terms should be zero under the mean-reversion hypothesis, we rerun these estimates without the constant term. The results are given in the right-hand part of the table. Then, the real USD/GBP exchange rate shows significant evidence of mean reversion on the 10% level when $T\hat{\delta}$ is used as the test statistic.

These results become somewhat sharper when the focus is changed from monthly to annual data, apparently because the use of annual data smooths over some monthly noise. Three of the real exchange rates then show significant signs of mean reversion without constant terms. However, the real USD/SEK rate continues to look off. When the constant is included, it is significantly positive on the 10% level. When it is not, neither test statistic shows anything close to significance of the δ coefficient. Although we believe that, fundamentally, the Swedish krona obeys purchasing power parity in the long run, we have to conclude that its performance on our sample does not qualify it as a candidate to illustrate the effects of mean reversion on that same sample. The three others, the Norwegian krone, the German mark/euro, and the British pound should, appear to pass this test, however.

In estimating annualized expected returns over varying horizons, we follow Mukherji (2011) by drawing 1000 independent samples of 120-month block returns with replacement from our data and use these samples to compute mean returns over horizons between one and ten years. We do this for the real returns of the respective portfolios valued in USD, NOK, DEM-EUR, and GBP. Because the USD was stronger at the

[3] This is certainly the case for the Swedish krona, the German mark (and the euro), as well as the British pound. Although the fixed exchange rate of the Norwegian krone was given up in December of 1992, the Norges Bank's mandate continued to require that the krone be kept "stable" against the other European currencies. This policy was effectively given up in mid-1998 and officially replaced by a formal inflation target in September of 2001.

[4] Data for S&P 500 total return were taken from Bloomberg, whereas the latter two indices were graciously made available to us by our colleague Espen Henriksen. The rates of return are defined as end-of-month to end-of-month log changes for monthly data and analogously for annual data. The exchange rates and CPIs were downloaded from the FRED database.

Table 1
Results of Dickey-Fuller tests.

	Monthly data						
	With constant terms				Without constant terms		
	Constant	$\widehat{\delta}$	$\widehat{\delta}/s.e.$	$T\widehat{\delta}$	$\widehat{\delta}$	$\widehat{\delta}/s.e.$	$T\widehat{\delta}$
S&P 500	0.010** (0.005)	−0.004 (0.005)	−0.840	−1.353			
FTSE GE	0.011** (0.006)	−0.008 (0.006)	−1.229	−2.574			
70–30 FTSE GE – Bloomberg GB	0.009** (0.004)	−0.007 (0.005)	−1.271	−2.324			
Real USD/NOK	0.002 (0.002)	−0.018 (0.012)	−1.541	−5.973	−0.013 (0.011)	−1.193	−4.316
Real USD/SEK	0.005* (0.003)	−0.018 (0.011)	−1.680	−5.310	0.002 (0.006)	−0.312	−0.631
Real USD/DEM-EUR	0.003 (0.002)	−0.020 (0.011)	−1.851	−6.699	−0.008 (0.008)	−0.998	−2.623
Real USD/GBP	0.001 (0.001)	−0.020 (0.012)	−1.633	−6.600	−0.020 (0.012)	−1.607	−6.507*
	Annual data						
	With constant terms				Without constant terms		
S&P 500	0.159* (0.087)	−0.088 (0.082)	−1.080	−2.291			
FTSE GE	0.154* (0.091)	−0.127 (0.105)	−1.205	−3.312			
70–30 FTSE GE – Bloomberg GB	0.116* (0.066)	−0.102 (0.088)	−1.159	−2.655			
Real USD/NOK	0.022 (0.023)	−0.270 (0.153)	−1.768*	−7.020	−0.218 (0.142)	−1.535	−5.678*
Real USD/SEK	0.059* (0.038)	0.216 (0.149)	−1.451	−5.611	−0.028 (0.088)	−0.318	−0.731
Real USD/DEM-EUR	0.038 (0.025)	−0.259 (0.144)	−1.800*	−6.639	−0.111 (0.109)	−1.012	−2.878
Real USD/GBP	0.004 (0.017)	−0.324 (0.165)	−1.966*	−8.427*	−0.326 (0.162)	−2.019**	−8.476*

Significance levels for $\widehat{\delta}/s.e.$ and $T\widehat{\delta}$ as tabulated by Fuller (1976).
Constant terms are t distributed.
Standard errors in parentheses.
* Significant at the 10% level.
** Significant at the 5% level.

end than at the beginning of our sample against all the three of the other currencies, the estimated mean returns are somewhat lower for the USD-valued portfolios. Under the maintained hypothesis of long-term purchasing parity, we ascribe these differences to sampling errors and adjust the returns of each of the non-USD portfolios at all horizons by subtracting the difference in the mean one-year return between the USD-based portfolio and the corresponding portfolios in the other respective currencies.[5]

The results are presented in Table 2 and illustrated in Fig. 2. Contrary to the evidence in Table 1, they indicate that the annualized expected returns decline with the horizon even in real U.S. dollars. This somewhat surprising result may reflect the well-known low power of the Dickey-Fuller test. However, the decline is clearly sharper for the real returns in Norwegian kroner or German marks/euros, as predicted by our theory. The decline does not start for real until after about four years and may even seem somewhat hump-shaped for the early years. However, for horizons beyond four years, the decline becomes substantial and proceeds approximately linearly; however, signs of the convexity seen in Fig. 1 can be detected between eight and ten years. We finally note that the GBP-based returns are very similar to the ones in U.S. dollars and even somewhat higher.

As suggested by the analysis in Section 4 above, the difference between the expected returns in USD and in the alternative currencies tends to be reduced by negative contemporary correlation between the respective exchange rates and the USD-based returns. If this correlation is strong enough it may erase the difference completely and even reverse it. Estimates of these correlations from our sample are presented in Table 3, with 95% confidence intervals based 10,000 bootstrap samples.

On monthly data, all the point estimates are significantly negative and of the same order of magnitude for all three currencies. They thus cannot account for the differences in performance between the GBP and the other two non-USD currencies shown in Table 2 and Fig. 2. However, the annual data show a different pattern, suggesting that the initial weakening effect of dollar-based financial losses may have been partly reversed in subsequent months for the NOK and the DEM-EUR. For the GBP, on the other hand, the difference goes in the opposite direction, so that the contemporaneous correlation is even more negative on annual than on monthly data. It is also significant or borderline significant.

This finding suggests that the special behavior of the expected returns on the real GBP-valued portfolios can indeed be explained by the negative contemporaneous correlation.

A remaining question is whether the equality in (14′) is satisfied for the real GBP-valued portfolios and thus explain the finding of slightly higher expected returns for this currency. Defining $\xi := \rho + \frac{1}{2}(\sigma_2/\sigma_1)$, this condition can be expressed equivalently for long horizons, i.e., for $t \to \infty$, as $\xi < 0$. Table 3 lists point estimates and 95% confidence intervals for this quantity as well. On monthly data, the point estimates are very close to zero and insignificantly different from zero in all cases. On annual data, the results again differ. For the real NOK and real DEM-EUR portfolios, the point estimates are all positive, indicating that (14′) is not

[5] The adjustments amounted to 0.90 pp. for NOK-based portfolios, 0.68 pp. for DEM-EUR, and 0.53 pp. for GBP.

Table 2
Adjusted expected annualized real rates of return (in percent) at varying horizons in different home currencies.

Portfolio	Horizon (years)	Currency			
		USD	NOK	DEM-EUR	GBP
S&P 500	1	7.17	7.17	7.17	7.17
	2	6.95	7.31	7.52	7.35
	3	7.35	7.30	7.52	7.41
	4	7.12	7.37	7.58	7.45
	5	6.28	6.46	6.68	6.65
	6	6.13	6.41	6.63	6.54
	7	5.58	5.53	5.81	5.86
	8	5.15	4.81	5.17	5.36
	9	4.83	4.17	4.60	4.89
	10	4.70	3.87	4.25	4.77
FTSE Global Equity	1	5.65	5.65	5.65	5.65
	2	4.92	5.27	5.48	5.31
	3	5.55	5.50	5.72	5.61
	4	5.24	5.48	5.69	5.57
	5	4.67	4.85	5.07	5.04
	6	4.57	4.85	5.07	4.98
	7	4.25	4.20	4.48	4.53
	8	3.95	3.61	3.97	4.16
	9	3.78	3.12	3.55	3.84
	10	3.79	2.97	3.34	3.86
	1	4.77	4.77	4.77	4.77
70–30 FTSE GE – Bloomberg GB	2	4.17	4.89	4.73	4.57
	3	4.69	4.79	4.87	4.75
	4	4.44	5.03	4.90	4.77
	5	4.02	4.45	4.42	4.40
	6	3.93	4.55	4.43	4.35
	7	3.73	4.08	3.97	4.02
	8	3.55	3.57	3.56	3.76
	9	3.47	3.20	3.24	3.53
	10	3.49	2.94	3.04	3.57

satisfied for $t \to \infty$. This result is consistent with the annualized expected values on these currencies being lower than the corresponding ones for the real USD-valued portfolios. For the real GBP-valued portfolios, the point estimates are all negative and rather large, suggesting that (14′) may indeed be satisfied for $t \to \infty$ for these portfolios, although not significantly so.

For shorter horizons, we define $\xi' := \rho + (\sigma_2/\sigma_1)$, so that the condition in (14′) can be expressed as $\xi' < 0$. The point estimates for this quantity (not shown) came out as positive in all cases, indicating that (14′) is not satisfied for short horizons.

6. Discussion

As an application of Corollary 1, the simulation study by Mork, Trønnes, and Bjerketvedt (2022) analyzes the future development of the Norwegian GPFG, whose investments are all made abroad in foreign-currency dominations, but whose returns are used to help defray government spending in Norwegian kroner. The estimation of their model on annual data found no significant contemporaneous correlation between the real USD equity returns and the rates of change in the USD/NOK real exchange rate, which is consistent with the findings in lower panel of Table 3.[6] Their findings imply an annual expected long-term real rate of return in domestic currency of close to 0.7 percentage points less than the one-period return. They also find that the expected rate of return falls quickly with a rising time horizon, such that the nearly 0.7 percentage point difference is essentially reached already at the five-year horizon, as illustrated in Fig. 1.

The results in Table 2 suggest a difference between the annualized ten-year and the one-year expected return of as much as 1.8 percentage points, more than twice the implications of the estimated simulation model. The rule for spending the returns of the GPFG stipulates annual withdrawals corresponding to the fund's expected short-term rate of return. As an application of Corollary 2, Mork et al. (op. cit.) estimate that the fund's expected real value 40 years hence is likely to be 18% lower than its current value of USD 1.3 trillion. This calculation ignores future deposits into the fund, which seem likely to be substantial given the current policies governing the Norwegian oil and gas sector. Even so, however, our findings throw considerable doubt on the sustainability of the current spending rule.

Out of the 1.8 percentage points of difference between ten-year and one-year annualized returns, 1.3 percentage points come from a corresponding difference between ten-year and one-year returns in real U.S. dollars. The latter point should be of considerable interest for the owners of U.S. endowment funds. The same would be true for U.K. funds, whose returns appear to follow the ones in U.S. dollars fairly closely. For funds based in Germany, our results fall in between the ones for British and Norwegian investors.

Dybvig and Qin (2021) go beyond expected values and show that a fund with a withdrawal policy like that of the GPFG eventually will lead to depletion even without serially correlation in the rates of return. The explanation comes from the same convexity as above, namely, that the distribution of the future fund value will become increasingly skewed to the right until, eventually, all probability mass is concentrated at a spike next to zero. Our analysis takes a step further by showing that, with negative serial correlation, the fund's future value cannot even be preserved in expectation, unless, that is, the instantaneous withdrawal rate is lowered from μ to $\mu - \sigma_2^2/2$.

Negatively serially correlated rates of return are not all bad news, of course, because they also mean that the variance rises more slowly than the time horizon, as seen in (9). The literature on mean reversion often refers to the variance ratio, defined as the ratio between the annualized t-period variance and the one-period variance. In our model, it becomes

$$\mathbb{VR}_t = \frac{\sigma_1^2 + \sigma_2^2 f(2\alpha t) + 2\rho\sigma_1\sigma_2 f(\alpha t)}{\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2}. \tag{16}$$

As expected, this ratio is less than unity if $\rho = 0$. This is the implicit assumption underlying the studies by, for example, Poterba and Summers (1988), Mukherji (2011), or Pástor and Stambaugh (2012). For the GPFG, under the assumption of $\rho = 0$, this ratio works out as 0.37 after 40 years, according to the findings by Mork et al. (op. cit.).

For $\sigma_2 \le \sigma_1$, the partial derivative of $\mathbb{VR}_t$ with respect to ρ is negative and larger in absolute value the larger αt. Thus, ceteris paribus, a negative contemporaneous correlation makes $\mathbb{VR}_t$ larger and may indeed cause it to exceed unity. It is easily seen that the condition for the latter to happen is identical to the condition that the annualized expected rate of return exceed the instantaneous one. Thus, if $\bar{r}_{0t} > \mu$, we should have $\mathbb{VR}_t > 1$, and vice versa.

In an apparent contradiction of our model, Mukherji (op. cit.) reports expected returns that rise or stay constant with rising horizons even in cases where the variance ratio falls well below unity. In our sample, we also seem to observe some elements of a similar contradiction between the indications of mean reversion from the variance ratios (not shown) and the expected annual returns. Further investigation of this apparent puzzle goes beyond the scope of the current paper, however.

7. Conclusion

Compounding introduces an element of convexity in the computation of long-term returns.

The implication that the distribution of long-term returns tends to be skewed to the right has been duly noted and analyzed in the literature. However, as long as the period rates of return are serially uncorrelated, the annualized expected long-term rate of return equals the short-term

[6] Mork et al. (op. cit.) did find a significant correlation of −0.44 between the real USD return of the bond part of the portfolio and the log level of the real exchange rate.

Fig. 2. Adjusted expected annualized real returns at varying horizons in different home currencies.

one. In this paper, we show how this implication unravels in the presence of negative serial correlation of returns.

Negative serial correlation has previously been studied for stock-market returns, with the effect on long-term risk as its main focus. This paper has moved the focus to expected returns and warns that long-term returns tend to be lower than short-term returns in the presence of negative serial correlation. Our main application concerns investors in small countries that invest their funds in foreign securities denominated in foreign currencies, but whose interest lies in their real rates of return in the domestic currency. Then, if the exchange rate obeys long-term purchasing power parity, the real rates of return in the domestic currency will be negatively serially correlated, and the annualized expected

Table 3
Contemporaneous cross correlations between the real USD returns and the rates of change of the respective real exchange rates.

		NOK	DEM-EUR	GBP
	Monthly data			
S&P 500	$\hat{\rho}$	−0.316	−0.235	−0.229
	95% conf. Interval	(−0.433, −0.189)	(−0.357, −0.110)	(−0.348, −0.108)
	$\hat{\xi}$	0.049	0.075	0.058
	95% conf. Interval	(−0.073, 0.183)	(−0.050, 0.204)	(−0.071, 0.183)
FTSE Global Equity	$\hat{\rho}$	−0.423	−0.343	−0.344
	95% conf. Interval	(−0.527, −0.307)	(−0.457, −0.222)	(−0.453, −0.228)
	$\hat{\xi}$	−0.062	−0.037	−0.062
	95% conf. Interval	(−0.175, 0.063)	(−0.050, 0.204)	(−0.180, 0.059)
70–30 FTSE GE – Bloomberg GB	$\hat{\rho}$	−0.482	−0.423	−0.386
	95% conf. Interval	(−0.578, −0.374)	(−0.527, −0.310)	(−0.491, −0.274)
	$\hat{\xi}$	0.009	−0.007	−0.002
	95% conf. Interval	(−0.107, 0.135)	(−0.115, 0.114)	(−0.119, 0.124)
	Annual data			
S&P 500	$\hat{\rho}$	−0.102	−0.054	−0.478
	95% conf. Interval	(−0.607, 0.498)	(−0.655, 0.377)	(−0.804, 0.067)
	$\hat{\xi}$	0.216	0.222	−0.216
	95% conf. Interval	(−0.344, 0.879)	(−0.134, 0.640)	(−0.550, 0.315)
FTSE Global Equity	$\hat{\rho}$	−0.223	−0.146	−0.559
	95% conf. Interval	(−0.655, 0.377)	(−0.502, 0.214)	(−0.839, −0.029)
	$\hat{\xi}$	0.089	0.125	−0.302
	95% conf. Interval	(−0.408, 0.776)	(−0.208, 0.544)	(−0.596, 0.226)
70–30 FTSE GE – Bloomberg GB	$\hat{\rho}$	−0.279	−0.229	−0.594
	95% conf. Interval	(−0.655, 0.377)	(−0.559, 0.113)	(−0.839, −0.029)
	$\hat{\xi}$	0.159	0.152	−0.223
	95% conf. Interval	(−0.335, 0.844)	(−0.162, 0.584)	(−0.596, 0.226)

long-term returns are lower than their short-term counterparts.

We find empirical support for this result on data for expected real returns at varying horizons for investors in Norway and Germany holding U.S. dollar denominated portfolios. Interestingly, a major part of the difference between long-horizon and short-horizon expected returns comes from a similar pattern for the same portfolios evaluated in real U. S. dollars. However, the effects of the currency exchange come on top of that. For U.K. investors we find no significant difference between returns in real British pounds or real U.S. dollars.

The difference between the German and Norwegian cases on the one hand and the British on the other, seems to arise from the fact the real rates of change in the USD/GBP exchange rate show significant negative contemporaneous correlation with the dollar-evaluated real portfolio returns. Such correlation dampens the effects of the negative serial correlation of the changes in the real exchange rates and may reverse it if it is sufficiently strong. The GBP case seems to lie on the borderline for this to happen.

We feel sovereign wealth funds, endowment funds, and other funds that aim for sustainability of returns would be well advised to take these results into account, especially when decisions are made about the allowable size of annual withdrawals. If specified as a percentage of the fund value, this percentage should in many cases be smaller than the rate of short-run expected returns.

Appendix A. Proof of Eq. (1)

Proof by induction:
Suppose the result holds for $n-1$, so that

$$\left(1+\overline{r}_{0,n-1}\right)^{n-1} := \mathbb{E}[(1+r_1)\cdots(1+r_{n-1})] = (1+\overline{r})^{n-1} + (1+\overline{r})^{n-3}\left[\sum_{j=1}^{n-2}(n-1-j)s_j + \text{h.o.m.}\right]$$

Then, using the rules for covariances of products (e.g. Bohrnstedt & Goldberger, 1969), we obtain

$$(1+\overline{r}_{0n})^n \equiv \mathbb{E}[(1+r_1)\cdots(1+r_n)] = \mathbb{E}\{[(1+r_1)\cdots(1+r_{n-1})](1+r_n)\}$$

$$= (1+\overline{r})\mathbb{E}[(1+r_1)\cdots(1+r_{n-1})] + (1+\overline{r})^{n-2}\left[\sum_{j=1}^{n-1}s_j + \text{h.o.m.}\right]$$

$$= (1+\overline{r})^n + (1+\overline{r})^{n-2}\left[\sum_{j=1}^{n-2}(n-1-j)s_j + \text{h.o.m.}\right] + (1+\overline{r})^{n-2}\left[\sum_{j=1}^{n-1}s_j + \text{h.o.m.}\right]$$

$$= (1+\overline{r})^n + (1+\overline{r})^{n-2}\left[\sum_{j=1}^{n-1}(n-j)s_j + \text{h.o.m.}\right],$$

which directly implies Eq. (1). For $n=2$, we see directly that

$$\mathbb{E}[(1+r_1)(1+r_2)] = (1+\bar{r})^2 + s_1.$$

Q.E.D.

Appendix B. Serial correlation of rates of return

Serial correlation comes only from the third term of (2). For the Ornstein-Uhlenbeck process, we know that the serial correlation of levels is given by the unconditional covariance

$$cov(x_t, x_s) = \frac{\sigma_2^2}{2\alpha} e^{-\alpha|t-s|}.$$

Positive serial correlation in levels implies negative serial correlation in rates of change. However, the correlation between the rates of change is complicated by the fact that they have order of magnitude *dt*. To circumvent this complication, we look first at finite differences:

$$cov(x_{t+h} - x_t, x_{s+h} - x_s)$$

$$= cov(x_{t+h}, x_{s+h}) - cov(x_t, x_{s+h}) - cov(x_{t+h}, x_s) + cov(x_t, x_s).$$

Without loss of generality, assume $h > 0$. Also, for simplicity, assume $t > s$; the case of $t < s$ is analogous. From the formula above, we then obtain

$$cov(x_{t+h} - x_t, x_{s+h} - x_s)$$

$$= \frac{\sigma_2^2}{2\alpha}\left(2e^{-\alpha(t-s)} - e^{-\alpha(t-s-h)} - e^{-\alpha(t-s+h)}\right)$$

$$= \frac{\sigma_2^2}{\alpha} e^{-\alpha(t-s)}\left[1 - \frac{1}{2}\left(e^{\alpha h} + e^{-\alpha h}\right)\right].$$

Because the Brownian motion part of the rate of return (2) is serially uncorrelated, this formula also describes the covariance between the rates of return r_t and r_s.

Similarly, the variance of the finite changes of the Ornstein-Uhlenbeck levels is

$$\mathbb{V}(x_{t+h} - x_t) = \frac{\sigma_2^2}{\alpha}\left(1 - e^{-\alpha h}\right).$$

To get the variance of the full rate of return we must, according to (2), add the variance of the Brownian motion part. For finite rather than infinitesimal asset-value changes, it can be written as

$$\mathbb{V}(r_t) = \sigma_1^2 h + \frac{\sigma_1^2}{\alpha}\left(1 - e^{-\alpha h}\right),$$

so that, for finite changes of the asset value,

$$corr(r_t, r_s) = e^{-\alpha(t-s)} \frac{1 - \frac{1}{2}\left(e^{\alpha h} + e^{-\alpha h}\right)}{\alpha h \sigma_1^2/\sigma_2^2 + 1 - e^{-\alpha h}}.$$

This expression becomes "zero over zero" for $h \to 0$. However, for small h,

$$\frac{\partial\left[1 - \frac{1}{2}\left(e^{\alpha h} + e^{-\alpha h}\right)\right]}{\partial h} = -\frac{1}{2}\alpha\left(e^{\alpha h} - e^{-\alpha h}\right) \to 0,$$

when $h \to 0$, and

$$\frac{\partial^2\left[1 - \frac{1}{2}\left(e^{\alpha h} + e^{-\alpha h}\right)\right]}{\partial h^2} = -\frac{1}{2}\alpha^2\left(e^{\alpha h} + e^{-\alpha h}\right) \to -\alpha^2.$$

Similarly,

$$\frac{\partial\left[\alpha h \sigma_1^2/\sigma_2^2 + 1 - e^{-\alpha h}\right]}{\partial h} = \alpha\left(\sigma_1^2/\sigma_2^2 + e^{-\alpha h}\right) \to \alpha\left(\sigma_1^2/\sigma_2^2 + 1\right),$$

and

$$\frac{\partial^2\left[\alpha h \sigma_1^2/\sigma_2^2 + 1 - e^{-\alpha h}\right]}{\partial h^2} = -\alpha^2 e^{-\alpha h} \to -\alpha^2$$

so, for small h,

$$corr(r_t, r_s) \approx -(1/2)e^{-\alpha(t-s)} \frac{\alpha^2 h^2}{\alpha(\sigma_1^2/\sigma_2^2 + 1)h - (1/2)\alpha^2 h^2}$$

$$= -\frac{\alpha}{2} e^{-\alpha(t-s)} \frac{h}{\sigma_1^2/\sigma_2^2 + 1 - (1/2)\alpha h} \approx -\frac{\alpha}{2} e^{-\alpha(t-s)} \frac{h}{\sigma_1^2/\sigma_2^2 + 1}.$$

Because the case where $s > t$ is analogous, we may thus write, for all infinitesimal changes

$$corr(r_t, r_s) = -\frac{\alpha}{2}e^{-\alpha|t-s|}\left(\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}\right)dt.$$

References

Arditti, F. D., & Levy, H. (1975). Portfolio efficiency analysis in three moments: The multiperiod case. *Journal of Finance, 30*(3), 797–809.

Bessembinder, H. (2018). Do stocks outperform treasury bills? *Journal of Financial Economics, 129*(3), 440–457.

Bohrnstedt, G. W., & Goldberger, A. S. (1969). On the exact covariance of products of random variables. *Journal of the American Statistical Association, 64*(328), 1439–1442.

Campbell, J. Y., & Viceira, L. M. (2002). *Strategic asset allocation: Portfolio choice for long-term investors*. Oxford: Oxford University Press.

Cho, J.-W., Choi, J. H., Kim, T., & Kim, W. (2016). Flight-to-quality and correlation between currency and stock returns. *Journal of Banking and Finance, 62*, 191–212.

Chortareas, G., & Kapetanios, G. (2009). Getting PPP right: Identifying mean-reverting real exchange rates in panels. *Journal of Banking & Finance, 33*(2), 390–404.

Cochrane, J. H. (2022). Portfolios for long-term investors. *Review of Finance*, 1–42.

Dybvig, P. H., & Qin, Z. (2021). *How to squander your endowment: Pitfalls and remedies*. Available at SSRN: https://ssrn.com/abstract=3939984. or.

Fama, E. F., & French, K. R. (1988). Permanent and temporary components of stock prices. *Journal of Political Economy, 96*, 246–273.

Fuller, W. A. (1976). *Introduction to statistical time series*. New York: John Wiley.

Hebisha, Z. (2023). Testing purchasing power parity in the Eurozone: A unit-root analysis. *MSA-Management Science Journal, 2*(2), 126–150.

Hossfeld, O., & MacDonald, R. (2015). Carry funding and safe have currencies: A threshold regression approach. *Journal of International Money and Finance, 59*, 185–202.

Maller, R., Müller, G., & Szimayer, A. (2009). Ornstein-Uhlenbeck processes and extensions. In T. Mikosch, J.-P. Kreiß, R. A. Davis, & T. G. Andersen (Eds.), *Handbook of financial time series*. Springer Verlag.

Mork, K. A., Trønnes, H. A., & Bjerketvedt, V. S. (2022). Capital preservation and current spending with sovereign wealth funds and endowment funds: A simulation study. *International Journal of Financial Studies, 10*(67), 1–24.

Mukherji, S. (2011). Are stock returns still mean-reverting? *Review of Financial Economics, 20*, 22–27.

Ning, C. (2010). Dependence structure between the equity market and the foreign exchange market—A copula approach. *Journal of International Money and Finance, 29*, 745–759.

Pástor, L., & Stambaugh, R. F. (2012). Are stocks really less volatile in the long run? *Journal of Finance, 58*(2), 431–477.

Poterba, J. M., & Summers, L. H. (1988). Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics, 22*, 27–59.

Rabe, C., & Waddle, A. (2020). The evolution of purchasing power parity. *Journal of International Money and Finance, 109*, Article 102237.

Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the theory of the Brownian motion. *Physical Review, 36*, 823–841.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics, 5*, 177–188.