NORGES BANK
INVESTMENT MANAGEMENT

01|2022

EXPECTED RETURNS ON A MULTI-ASSET PORTFOLIO

DISCUSSION NOTE

We outline a framework for estimating expected returns on equities, government bonds, and corporate bonds. We combine estimates of expected returns across asset classes in developed markets to obtain the expected return on a multi-asset portfolio that approximates the benchmark index of the Government Pension Fund Global (GPFG).

Date 19/12/2022
ISSN 1893-966X

dn@nbim.no
www.nbim.no

SUMMARY

- We outline a framework for estimating expected returns on equities, government bonds, and corporate bonds, which we apply to assets in the US, the euro area, Japan, and the UK. Across asset classes, we estimate expected returns as the sum of three components: income, cash flow growth and valuation changes.
- Our approach emphasises the use of forward-looking data, combining market-implied expectations from traded assets and survey-based forecasts. This stands in contrast to other commonly used approaches, which tend to be based on historical returns and realised fundamentals. Forward-looking estimates of expected returns perform better than estimates based on historical returns, which can be biased and excessively volatile.
- Focusing on long-horizon expected returns, our estimates indicate that expected returns were declining across asset classes between the Global Financial Crisis (2007-2009) and the outbreak of the Covid-19 pandemic. Since then, long-term expected returns have been increasing. As at the end of Q3 2022, expected real returns on equities and government bonds are estimated to be around 3.8 and 0.7 percent, respectively. Declining expected returns have predominantly been driven by falling real interest rates across developed markets. Risk premiums associated with exposure to equity and corporate bond markets have slightly increased and thus partially offset the decline driven by real rates.
- We combine estimates across asset classes to obtain the expected return on a portfolio that approximates the benchmark index of the Government Pension Fund Global (GPFG). The expected real return on this multi-asset portfolio has been declining over the last decade, with part of the decline reversing in the post-Covid-19 period, reaching 3 percent at the end of Q3 2022. It should be emphasised that our estimates are uncertain and can vary substantially over time.

1. Introduction

Estimates of expected returns are a key input into any investment decision. For investors holding diversified multi-asset portfolios, expected return estimates need to span most major asset classes and markets across a range of investment horizons.

In this note, we outline a framework for estimating expected returns on three major asset classes: equities, government bonds, and corporate bonds. For each asset class, we estimate expected returns as the sum of three components: income, cash flow growth and valuation changes. We apply our framework to equity and bond markets in the US, the euro area, Japan, and the UK. We apply the methodology to key asset classes in developed markets with the aim of producing expected returns for a portfolio similar to the benchmark index of the Government Pension Fund Global (GPFG). While we focus on expected returns over long horizons, our approach allows us to estimate expected returns at shorter horizons as well.

Our approach emphasises the use of forward-looking data, combining market-implied expectations from traded assets and survey-based forecasts. This modelling choice separates our framework from other commonly used approaches. Most of these approaches rely either directly or indirectly on historical returns and fundamental metrics.

Our estimates of expected equity returns have been declining over the past decade. This decline has been driven by decreasing cash flow growth expectations and a compression in dividend yields. As at the end of Q3 2022, the long-term expected real return on equities is estimated to be approximately 3.8 percent. The expected real return on government bonds was also falling between the Global Financial Crisis (GFC) and the outbreak of the Covid-19 pandemic. Since then, the expected return on government bonds has been increasing and is estimated to be around 0.7 percent at the end of Q3 2022. We show that steadily declining bond yields have been the main driver of lower expected returns across both asset classes.

We combine estimates across asset classes to estimate the expected real return on a multi-asset portfolio. This estimate has been declining over the last decade, with part of the decline reversing in the post-Covid-19 period, reaching 3 percent at the end of Q3 2022.

Compared to alternative estimates based on historical returns, our approach produces expected equity return estimates that differ significantly in terms of both their level and variation over time. These differences are predominantly driven by the cash flow growth and valuation components of expected returns. Alternative approaches based on backward-looking data produce estimates of expected cash flow growth and valuation changes that vary considerably more than those estimated using our forward-looking framework. The volatile nature of backward-looking estimates would imply a level of equity volatility that is hard to reconcile with observed returns.

The note proceeds as follows. In the next section, we motivate the use of forward-looking data in the estimation of expected returns. Section 3 outlines the expected return framework by asset class. In Section 4, we present our estimates of expected returns for each asset class and for multi-asset portfolios. Section 5 compares our estimates to backward-looking approaches to estimating expected returns. Section 6 concludes.

2. Forward-looking Estimation of Expected Returns

Realised asset returns can be broken down into two components representing *expected* and *unexpected* returns:

$$R_{t+1} = E_t\left(R_{t+1}\right) + e_{t+1}, \tag{1}$$

where R_{t+1} is the return realised between t and $t+1$, $E_t(R_{t+1})$ is the expected return based on information available at time t, and e_{t+1} is the unexpected return. The unexpected return refers to any unanticipated changes in asset prices that cause a wedge between expected returns $E_t(R_{t+1})$ and realised returns R_{t+1}.

Our goal is to obtain a precise estimate of the conditional expected return, $E_t(R_{t+1})$. This goal is more ambitious than estimating the unconditional expected return $E(R_{t+1})$, which can be obtained as a sample average of realised returns.[1] Equation (1) suggests that, in principle, it should be possible to extract the conditional expected return from realised returns. A simple way of extracting expected returns from realised returns is to estimate a moving average of returns. The idea behind applying a moving average is that it smooths through unexpected return variation and thus extracts the expected return. In setting the size of the rolling window, investors face a trade-off between smoothing through unexpected returns, which requires a longer window, and the timeliness of the estimate, which requires a shorter window.

Estimates based on moving averages of past returns rely on the assumption that the cumulative effect of unexpected returns is negligible. However, if unexpected returns are large in magnitude or correlated over time, their cumulative effect on return-based estimates of expected returns can be sizeable and persistent.[2] This turns out to be the case for both equity and bond returns.

For illustration, we apply equation (1) to government bond returns and write the realised log return on a zero-coupon government bond with duration of n periods as follows:

$$r_{t+1}^{(n)} = \underbrace{y_t^{(n)}}_{\approx \text{ expected return}} + \underbrace{(n-1)\left(y_t^{(n)} - y_{t+1}^{(n-1)}\right)}_{\approx \text{ unexpected return}}. \tag{2}$$

---

[1] It is well known that sample averages of realised returns are imprecise estimates of unconditional expected returns even in sample periods that span several decades.

[2] Indeed, empirical evidence suggests that unexpected returns drive a significant portion of the volatility of realised returns, see e.g. Merton (1980); Elton (1999). This is especially the case for long-duration assets such as equities or long-term bonds. Empirical evidence also suggests that shocks to expected returns are negatively correlated with unexpected returns. Pástor and Stambaugh (2009) provide a thorough theoretical and empirical discussion of this phenomenon.

Equation (2) shows that the volatility of unexpected returns increases with duration. As a consequence, realised returns on long-duration bonds are dominated by unexpected returns. The steady decline in government bond yields over the last four decades has been accompanied by positive unexpected returns that have been particularly large for long-duration bonds. Unexpected returns like these are unlikely to be repeated in the future.

Equity returns also exhibit the yield dynamics outlined above. In addition, rare but large shocks to dividend growth expectations, and their corresponding risk premiums, are important determinants of realised equity returns. Recent examples of such shocks include the GFC and the market drawdown associated with the outbreak of the Covid-19 pandemic.

Despite these issues, most expected return estimates rely on historical data in some form. In an extreme case, expected return estimates based on historical returns that do not account for the properties of unexpected returns can be completely unrelated to the future returns they are meant to predict.[3]

A key modelling choice in our approach is therefore either to avoid using realised returns in the estimation of expected returns, or to explicitly account for the issues outlined above when doing so. Our estimates of expected returns rely on forward-looking data, predominantly a combination of market-implied expectations and survey-based forecasts. This is the most important modelling choice that separates our implementation from other approaches. We explore this issue in more detail by comparing our estimates of expected returns with alternative estimates based on historical data in Section 5.

3. Defining Expected Returns by Asset Class

In this section, we outline our framework for estimating expected returns across asset classes. Our expected return framework is closely related to the sum-of-the parts (SOP) framework. The general SOP framework is a common starting point and is well-documented in the literature.[4] The central premise of the SOP framework is to break asset returns into three components – income, cash flow growth, and valuation changes, commonly referred to as "carry", "growth", and "value", respectively. Each component is then estimated separately, and the total return expectation is the sum of the three components. This approach avoids many of the issues associated with earlier approaches to return forecasting. Most of these earlier approaches are essentially based on freely estimated predictive relationships, commonly obtained by regressing total returns on various forecasting variables. Expected return estimates based on these approaches, however, have been found to perform poorly when evaluated out-of-sample. This is predominantly due to over-fitting in a given sample period, resulting in

---

[3] Persistent variation in asset returns drives a wedge between unconditional and conditional asset return expectations. It not only makes it harder to estimate the unconditional expected returns, but also makes them less relevant relative to the conditional expectations. Throughout the note, we focus on the conditional expectations about expected asset returns.

[4] See e.g. Blanchard (1993); Fama and French (2002); Ibbotson and Chen (2003); Ferreira and Santa-Clara (2011); Ilmanen (2011).

unstable and counter-intuitive predictive relationships.[5] The SOP framework largely avoids these pitfalls by imposing economic priors and estimating each of the three return components separately. Our SOP implementation is based on the same three return components, but we model each component using forward-looking data.

Our focus is on estimating expected returns over long investment horizons where valuation changes play only a minor role. As a result, our estimates of long-term expected returns rely predominantly on the carry and growth components of asset returns. We also show how our approach can be used to estimate expected returns at shorter horizons by modelling the value component. We start by describing the framework for estimating expected *nominal* returns for each asset class. Expected real returns are then obtained by subtracting survey-based inflation expectations from nominal estimates. Across asset classes, we focus on expected returns in local currency. Unless otherwise stated, we focus on asset markets in the US, the euro area, the UK and Japan, which we collectively refer to as the "G4" markets. Throughout the note, we work with annualised log-returns.

Expected Equity Returns

The static Gordon growth model is a common starting point for motivating the underlying components of expected equity returns. The model relates the price of an equity index to expected returns and dividend growth in the following way:

$$P_t = \frac{E(D_{t+1})}{E(R) - E(G)}, \tag{3}$$

where P_t is the current equity index price, $E(D_{t+1})$ is the expected value of next year's dividend, $E(R)$ is the discount rate, or equivalently the *expected return* that investors require to hold the equity index. Finally, $E(G)$ refers to expected dividend growth.

We rearrange equation (3) in terms of expected equity returns:

$$E(R) = \underbrace{\frac{E(D_{t+1})}{P_t}}_{\text{Carry}} + \underbrace{E(G)}_{\text{Growth}}. \tag{4}$$

In this model, the expected equity return is constant, and its level is determined by two components: the dividend yield (carry) and expected dividend growth (growth). The carry component tells us how much an investor would expect to earn on their equity holdings if the index dividend stayed at its current level throughout the holding period. The growth component represents the contribution from future dividend growth.

The static Gordon growth formula assumes constant dividend growth $E(G)$. This assumption, however, is not very realistic as expected dividend growth varies over time. Relaxing the assumption of constant dividend growth

---

[5] See Welch and Goyal (2008).

results in a dynamic version of equation (3):

$$P_t = \sum_{i=1}^{\infty} \frac{E_t D_{t+i}}{\left(1 + R_{t,\infty}\right)^i}. \tag{5}$$

Linearising equation (5) and solving for $R_{t,\infty}$, which we refer to as "expected return annuity", we get:[6]

$$R_{t,\infty} \approx \underbrace{E_t\left(D_{t+1}\right)/P_t}_{\text{Carry}} + \underbrace{G_{t,\infty}}_{\text{Growth}}. \tag{6}$$

The second term in equation (6), which we refer to as the "growth annuity", is a weighted sum of dividend growth rates across all periods, given by:

$$G_{t,\infty} = \sum_{n=2}^{\infty} w_t^{(n-1)} E_t \Delta d_{t+n}, \tag{7}$$

where Δd_{t+n} refers to nominal dividend growth between $t+n-1$ and $t+n$. The weights $w_t^{(n)}$ in equation (7) are determined as the ratio of the present value of the dividend payout at time $t+n$ and the present value of all future dividends, which corresponds to the current equity index price P_t. For long-duration assets such as equities, short-term cash flow growth expectations receive a relatively small weight, making their long-horizon counterparts key determinants of the growth annuity.

The expected return annuity $R_{t,\infty}$ can be separated into a risk-free component R_t^f[7] and the equity risk premium ERP_t:

$$R_{t,\infty} = R_t^f + ERP_t. \tag{8}$$

This allows us to rewrite equation (6) as follows:

$$R_t^f + ERP_t \approx E_t\left(D_{t+1}\right)/P_t + G_{t,\infty}. \tag{9}$$

Equation (9) states that the expected equity return can be expressed in two different ways, either as a sum of the equity risk premium and the risk-free rate, or by combining the dividend yield with the growth annuity.

The expected return annuity can be interpreted as the internal rate of return that investors would earn if they held the equity index forever. For such buy-and-hold investors, the expected return depends on the current stock price and the sequence of expected dividends.[8] While an infinite investment horizon might be unrealistic, these estimates serve as a good approximation of expected returns at sufficiently long horizons, e.g. ten years or longer.

If an investor's horizon is shorter than buy-and-hold for equities, predictable variation in expected cash flows and discount rates plays a role in determining

---

[6] Appendix A provides a more detailed derivation of the formulas given by equations (6) and (7).

[7] Here, we refer to a generic risk-free rate without specifying its maturity. We make the definition more precise later in the note.

[8] After the initial purchase, price changes matter for buy-and-hold investors only to the extent they are driven by changing dividend growth expectations (cash flow risk). Price changes driven by shifts in risk premiums or interest rates (discount rates) are irrelevant for buy-and-hold investors.

expected returns. This variation introduces an additional component alongside carry and growth, which we label "value". To be parsimonious in our implementation, we assume that time-varying risk premiums are the dominant source of predictable variation in expected returns at shorter horizons.[9]

The formula for long-horizon expected equity returns in equation (6) then becomes:

$$E_t\left(R_{t,h}\right) \approx \underbrace{E_t\left(D_{t+1}\right)/P_t}_{\text{Carry}} + \underbrace{E_t\left(G_{t,\infty}\right)}_{\text{Growth}} \underbrace{-E_t\left(\Delta ERP_{t+h}\right)}_{\text{Value}}, \tag{10}$$

where ΔERP_{t+h} is the expected change in the equity risk premium over horizon h.

To estimate the long-horizon expected equity return, we need an estimate of the term structure of dividend growth and the corresponding weights. We estimate term structures of expected dividend growth and their corresponding weights in equation (7) using the two-stage present-value model outlined in NBIM (2021b). Specifically, we express the equity index level as the sum of expected dividends discounted by the risk-free interest rate and the equity risk premium at a corresponding horizon. In the first stage, spanning the first 30 years, we model each annual index dividend explicitly. The second stage refers to all dividends beyond the 30-year horizon, which are represented through a perpetuity. The weights up to the 30-year maturity are obtained as a ratio of the dividend strips and the equity index, with the perpetuity receiving the remaining weight.

In the estimation, we apply a present-value restriction to ensure that the modelled equity index is as close as possible to the observed equity index at each point in time. The goal is to obtain estimates of expected dividend growth and their corresponding weights that are consistent with the market pricing.

The carry component is approximated by the expected index dividend, 1-year ahead, divided by the current equity index level. For the growth component, we need to weight estimates of expected dividend growth across the entire term structure to get an estimate of the growth annuity. In many cases, the dividend perpetuity beyond 30 years represents half or more of the total present value of expected dividends.

To estimate the value component, we construct the risk premium annuity using estimates of the term structure of dividend risk premiums presented in NBIM (2021b). We use the same weights as for the growth annuity. We then compare the risk premium annuity to a 10-year rolling average of its past values. The value component is set equal to the difference between the latest estimate of the risk premium annuity and the rolling average.

The present-value model in NBIM (2021b) relies on the use of dividend futures, prices of which are informative of the market value of dividends at

---

[9] It is common to express the value component as an expected change in the valuation ratios, see e.g. Ferreira and Santa-Clara (2011).

fixed horizons stretching several years into the future. We combine these futures prices with stock index prices, option data and surveys within the present-value framework. Specifically, we use futures data from Bloomberg and equity index options data from OptionMetrics as our main data sources. Equity index price levels are sourced from Bloomberg. We use survey-based estimates of nominal GDP growth at the ten-year horizon from Consensus Economics to identify the long end of expected dividend growth. Finally, we use inflation forecasts across all available horizons from Consensus Economics to estimate the expected real returns across G4 markets. Due to a relatively short history of data available on dividend futures, we can only extend the sample period back to January 2003.

Expected Government Bond Returns

The perspective of a buy-and-hold investor is also a common starting point when estimating long-term expected returns on government bonds. When doing so, the expected log return on a portfolio of government bonds is given by its initial yield $y_t^{(n)}$:

$$E_t\left(R_{fi,t+n}^{(n)}\right) = y_t^{(n)}, \tag{11}$$

where n refers to the duration of the portfolio. This simple relationship holds more generally. In fact, the return on any government bond portfolio with relatively stable duration will eventually converge towards the starting yield at sufficiently long horizons. The convergence happens as price returns are offset by accruals over time. For instance, if interest rates go up, bond prices adjust downwards, but future returns will be higher, such that the total return on the bond investment will be close to the initial yield.[10] This result naturally also applies to broadly diversified fixed income benchmark indices. We therefore use bond index yields to approximate expected nominal bond returns at long investment horizons.

At shorter investment horizons, bond yields and returns fluctuate around their long-term expected returns. Empirical evidence suggests that these short-term fluctuations are primarily driven by the so-called term premium (Bauer and Rudebusch, 2020; Feunou and Fontaine, 2021). The term premium is the compensation an investor receives for being exposed to duration risk, which refers to uncertainty around the future evolution of short rates. The premium also reflects yield moves driven by temporary fluctuations in aggregate demand for safety and liquidity, often referred to as the convenience yield.

We take short-term yield fluctuations into account when estimating expected returns at shorter investment horizons h, which we express as:

$$E_t\left(R_{fi,t+h}^{(n)}\right) \approx \underbrace{y_t^{(h)}}_{\text{Carry}} + \underbrace{E_t\left(RX_{t+h}^{(n)}\right)}_{\text{Value}}, \tag{12}$$

---

[10] For a detailed discussion on this topic, see e.g. Leibowitz and Bova (2012).

where $y_t^{(h)}$ denotes a zero-coupon yield with maturity h and $E_t\left(RX_{t+h}^{(n)}\right)$ is the expected excess return on an n-period bond, which we refer to as the "value" component. Note that $E_t\left(RX_{t+h}^{(n)}\right)$ refers to a specific investment horizon h and is thus related to but distinct from the term premium.[11]

The carry component for an h-period investment horizon can be represented by the yield on a government bond with a maturity of h periods. The total expected return on an n-period government bond at horizon h is the sum of the carry and the value components as indicated in equation (12).

Our estimation of expected excess returns needs to account for the persistent variation in bond returns. We estimate expected excess returns using a predictive regression estimated in real time. In the first step, we remove persistent variation from yields by subtracting long-horizon inflation expectations and the equilibrium real rate. The outputs from this step are yield "cycles" which contain the time-varying risk premium and the transitory part of short rate expectations. The second step involves regressing the realised bond excess return on a 10-year bond on the short- and long-maturity yield cycle. The fitted value from this regression is an estimate of the expected excess return on a 10-year bond. More details on this methodology are provided in Cieslak and Povala (2015).

We use the yield-to-maturity of country-level government bond indices from Bloomberg Fixed Income Indices. For maturity-specific yields, we source zero-coupon nominal spot yields from ICE Indices.[12] The data are available at a daily frequency and at three-month maturity steps between three months and 30 years. We start in January 2003 to align with other asset classes.

Expected Corporate Bond Returns

Similar to equities and government bonds, the perspective of a buy-and-hold investor is a common starting point when estimating expected returns on corporate bonds. From this perspective, the expected return on a corporate bond is represented by the carry component, which has two parts: (i) the expected return on a nominal government bond of the same duration as the corporate bond $y_t^{(n)}$ and (ii) the expected excess return on the corporate bond, CRP_t, usually referred to as the credit risk premium (CRP). The credit risk premium compensates the investor for being exposed to corporate credit risk. The expected return on a corporate bond with an n-period duration for a buy-and-hold-investor is given by:

$$E_t\left(R_{cr,t+n}^{(n)}\right) \approx y_t^{(n)} + CRP_t. \tag{13}$$

We can estimate expected returns on the first part of the carry component using the same methodology as outlined for government bonds. It is less appropriate, however, to apply the same approach to estimate the second

---

[11] The term premium refers to the average of the expected excess returns over the lifetime of a bond. Hence, $E_t\left(RX_{t+h}^{(n)}\right)$ and the term premium coincide when h is equal to the bond's maturity.

[12] Formerly Bank of America/Merrill Lynch Fixed Income Indices.

part. Excess returns on corporate bond indices do not converge to initial corporate spreads in the same way as in the case of government bonds. This is due to issuance-induced index dynamics, defaults and rating migration, particularly during periods of market stress, as argued in Desclée and Polbennikov (2018).

We use duration-matched government bond yields to represent the first component of expected returns, and focus on estimating the second component. This component is usually referred to as the credit risk premium (CRP) and is defined as follows:

$$CRP_t = CS_t - L_t, \tag{14}$$

where CS_t refers to the yield differential between corporate bonds and maturity-matched government bonds and L_t refers to the loss an investor expects to incur in case of default.[13]

Following this relationship, we estimate expected credit risk premiums for investment-grade corporate bonds denominated in USD and EUR. The overall expected excess return on corporate bonds is the market-value-weighted sum of the two market-specific estimates. Corporate bonds denominated in USD and EUR account for the majority of investment-grade corporate debt issuance, so by focusing on these two currencies we are able to capture most of the market dynamics.[14]

We estimate expected credit risk premiums by subtracting the product of the default probability and the loss given default from the credit spread. While this adjustment is sufficient for a buy-and-hold investor, there are other factors that may potentially impact future returns for investors that closely track a corporate bond index. These factors include, for instance, the so-called fallen-angel effect.[15] Most of these factors, however, are due to index dynamics or are specific to a given sample period, so we choose to leave them out of our estimates of long-term expected returns.

We obtain option-adjusted credit spreads from currency-level corporate bond indices from Bloomberg Fixed Income Indices. We follow the methodology in Culp, Nozawa, and Veronesi (2018) to estimate the probability of default. Specifically, we use a no-arbitrage condition between the risk-free debt and equity index option pricing to construct a portfolio of pseudo investment-grade corporate bonds, and estimate their probability of default. We combine these estimates with historical loss-given-default rates informed by Moody's corporate default and recovery rates study. We provide more details on how we estimate the expected credit loss in Appendix B.

---

[13] The expected credit loss is a product of the loss given default and the likelihood of default.

[14] For example, as at December 2021, USD- and EUR-denominated corporate bonds made up 67 percent and 24 percent of the Bloomberg Barclays Corporate Global Investment Grade Index, respectively.

[15] See Ilmanen (2011) for a comprehensive overview of the factors that have impacted historical returns for investors that track corporate bond indices.

4. Estimates of Expected Returns

In this section, we present estimates of expected returns across asset classes, and combine them to estimate the expected return on a multi-asset portfolio that approximates the benchmark index of the GPFG. We present estimates for a sample period ending in Q3 2022.

Equity Expected Return Estimates

Figure 1 shows expected returns for G4 equity markets. Each panel shows expected nominal returns alongside their long-term components: carry and growth. When comparing estimates across markets, US equities stand out in several ways. Expected cash flow growth has been the largest component of expected US equity returns throughout the sample period. This differs from the other markets, where its contribution towards overall expected returns is approximately on a par with the carry component. Despite a better growth outlook, overall US expected returns are lower than those in the euro area or the UK at the end of the sample. This is because the comparatively high growth component is already reflected in US equity prices, thus depressing the carry component. When comparing growth components across markets, the US components are not only higher, but also more stable over time.

Expected returns on US and euro area equity markets – the two largest markets among the G4 – have been declining for about a decade. This decline has been driven by a combination of gradually declining cash flow growth expectations and a compression in their equity carry components. Declining return expectations have coincided with increasing prices in both equity markets. While valuations do not influence the growth component, declining carry components are directly linked to higher equity prices.

In contrast to equity markets in the US and the euro area, expected return estimates for Japanese and UK equities have not been trending downwards in our sample period. Long-term expected returns on Japanese equities have been considerably lower than the other estimates. Since these are nominal returns, this can be partly attributed to comparatively low inflation expectations in Japan. Finally, while the initial phase of the Covid-19 pandemic had a large negative impact on carry components in the UK and the euro area, the impact was more muted across equity markets in Japan and the US.

Figure 2 shows the same expected return estimates as in Figure 1, alongside the value component estimates. While the value component is more relevant at shorter horizons, it tends to fluctuate around zero and make up a smaller share of total return expectations than the other two components.

Figure 3 shows aggregated expected returns on G4 equities, including nominal and real estimates. To obtain aggregate expected equity returns, we combine G4 estimates using market-cap weights. Both nominal and real expected equity returns have been steadily declining over the last decade. While the GFC is barely visible in these long-horizon expectations, the Covid-19 pandemic triggered a sharp drop in expected equity returns,

Figure 1: Estimates of long-term expected equity returns

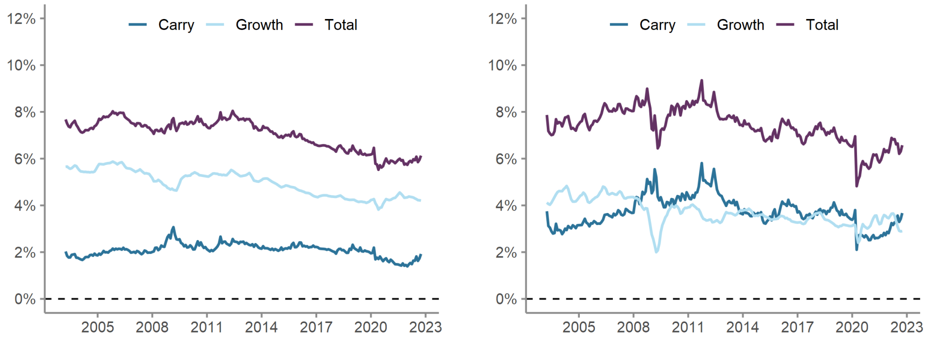

(a) United States *(b) Euro area*

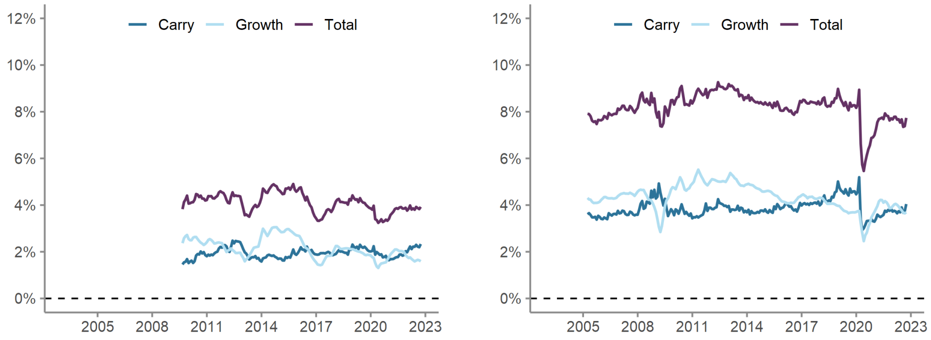

(c) Japan *(d) United Kingdom*

Note: The figure shows estimates of expected equity returns. The sample period ends in September 2022. The start of the sample period varies by market: January 2003 for the euro area and the US, July 2009 for Japan, and January 2005 for the UK.

predominantly driven by the carry component. As at the end of Q3 2022, the long-horizon expected real return on G4 equities is estimated to be approximately 3.8 percent.

Equity Risk Premium Estimates

The expected excess return of equities over government bonds – or equity risk premium (ERP) – is a measure of how much investors can expect to be compensated for being exposed to equity market risk. Our framework allows us to estimate a forward-looking version of the ERP. While most of the literature defines the ERP as the average excess equity return over short-maturity government bonds, we define the ERP as the expected equity return in excess of the expected return on a government bond with a duration that matches the equity duration:

$$ERP_t = R_{t,\infty} - y_t^m, \qquad (15)$$

Figure 2: Estimates of expected equity returns including the value component

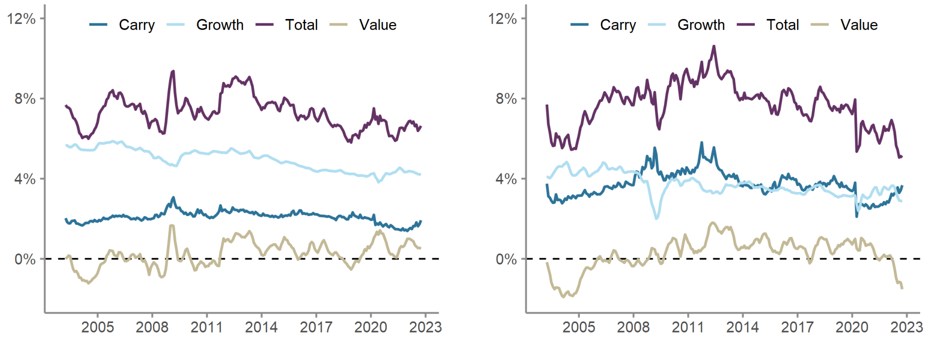
(a) United States

(b) Euro area

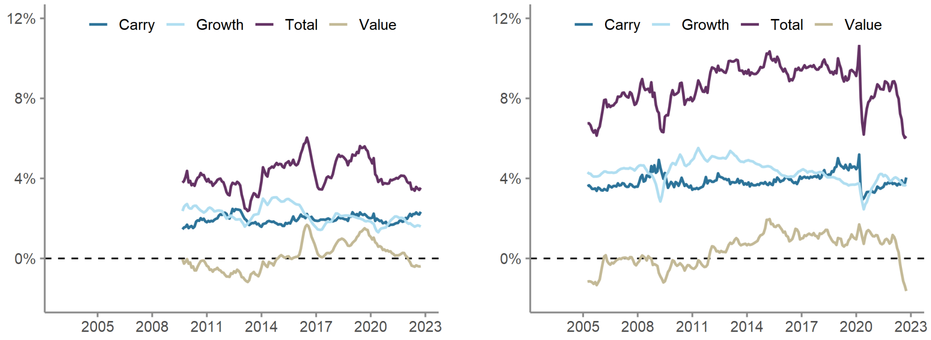
(c) Japan

(d) United Kingdom

Note: The figure shows estimates of expected equity returns. The sample period ends in September 2022. The start of the sample period varies by market: January 2003 for the euro area and the US, July 2009 for Japan, and January 2005 for the UK.

Figure 3: Aggregate long-term expected equity return estimate

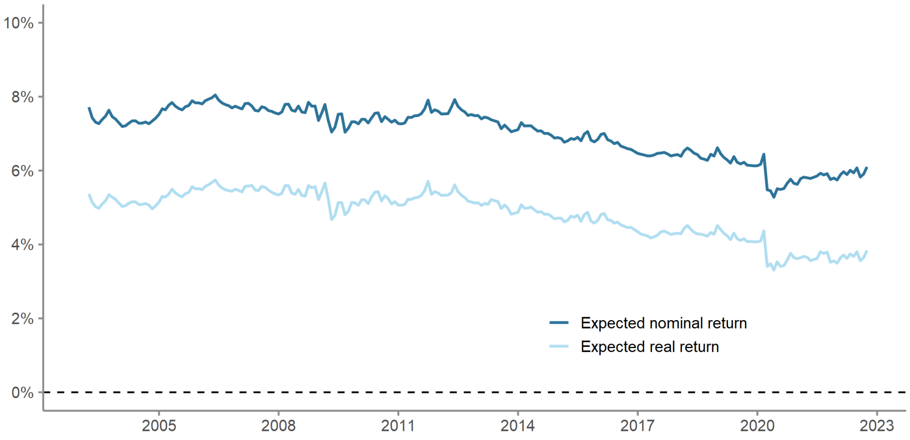

Note: The figure shows market-cap-weighted estimates of expected G4 equity returns. The expected real return is obtained by subtracting survey-based horizon-matched inflation forecasts from the expected nominal return. The sample period ends in September 2022.

where m refers to the equity duration.[16] This definition of ERP follows from equation (8) where we approximate the risk-free rate using a long-duration

---

[16] The model presented in NBIM (2021b) allows us to estimate the duration of equity cash flows at each point in time. The equity duration is a function of the maturities of the cash flow stream and

government bond yield.

Our definition of the ERP accounts for the fact that the appropriate risk-free asset depends on the investor's horizon. From the perspective of a long-term investor, long-duration bonds are risk-free assets, while short-duration bonds are risky.[17] To assess the attractiveness of equities relative to a risk-free investment, a long-horizon investor would use the ERP definition from equation (15).[18]

The declining return expectations shown in Figure 3 do not necessarily make equities a less attractive asset class, particularly when evaluated on a relative basis against government bonds. We can gauge the relative attractiveness of equities and government bonds using the forward-looking ERP defined in equation (15).

Figure 4 shows the evolution of these estimates for G4 equity markets. In contrast to the declining return expectations on G4 equities, the equity risk premium estimates have increased over time. Based on the relationships described in equations (9) and (15), this suggests that lower expected equity returns can be largely attributed to declining risk-free interest rates. Expected equity risk premiums are particularly high in the UK and euro area equity markets, both of which can be characterised by relatively volatile dividend streams and low dividend growth expectations, as was shown in Figure 1.[19]

The estimates are defined as expected equity returns in excess of the risk-free rate. While we define the risk-free rate as the yield on government bonds, with a duration that approximates the equity duration, a common alternative is to use the yield on short-term government bills to proxy the risk-free rate. This modelling choice only matters if yield changes are persistent, which is the case empirically. The duration of the risk-free asset can therefore have a meaningful impact on the estimates of the equity risk premium. van Binsbergen (2021) uses realised returns to show that this has indeed been the case historically. We compare our ERP estimate with a traditional measure of the ERP calculated using short-term bond yields in Figure 5. We use estimates of equity market duration from NBIM (2021b). The two ERP estimates are naturally positively correlated, as they are both based on the same estimate of expected equity returns. There is, however, a considerable wedge of several percentage points between the two series, which results from switching from short-term to duration-matched yields. The version of the ERP using short-term yields averages 5.9 percent over the sample period, while our ERP estimate has an average of 3.6 percent over the same period.

---

the corresponding weights $w_t^{(n)}$. For example, the estimates of duration for US equities range from 20 to 25 years.

[17] This is because investing in short-duration bonds and rolling over the investment when the bond matures exposes the investor to reinvestment risk. The relative importance of reinvestment risk is linked to the persistence of interest rates.

[18] A measure of the equity risk premium based on short-duration bonds does not account for the portion of equity returns that is generated by changes in long-term interest rates, which can be sizeable as documented in NBIM (2021a). Investors can gain exposure to this portion of equity returns by investing in long-duration government bonds.

[19] We explore the properties of expected dividend growth across horizons in NBIM (2021b).

Figure 4: Forward-looking estimates of the long-term equity risk premium

(a) United States *(b) Euro area*

(c) Japan *(d) United Kingdom*

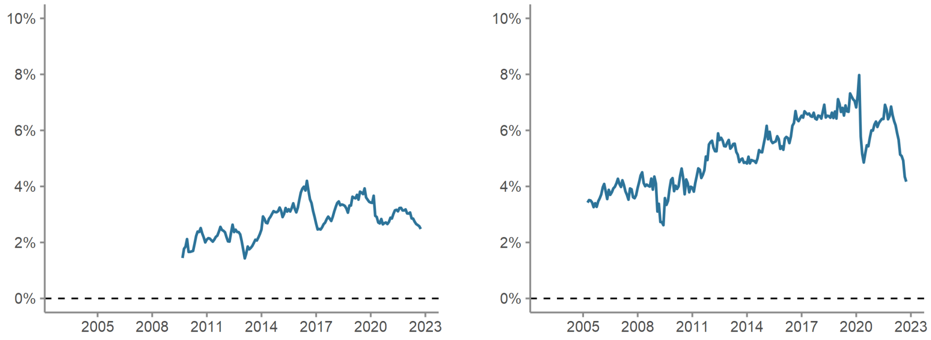

Note: The figure shows forward-looking estimates of the equity risk premium using long-term yields as a proxy for the risk-free rate, as defined in equation (15). *The sample period ends in September 2022. The start of the sample period varies by market: January 2003 for the euro area and the US, July 2009 for Japan, and January 2005 for the UK.*

Government Bond Expected Return Estimates

Figure 6 shows long-term estimates of expected nominal returns for G4 government bonds. Across bond markets, expected returns have been declining steadily over most of our sample period, eventually reaching zero in the euro area and in Japan. Recent months have seen a sharp reversal of the downward trend, driven by monetary policy tightening in response to high inflation. One notable exception is Japanese government bonds, with yields still being close to zero. Note that the expected return on Japanese government bonds reached zero around 2016. This coincides with the introduction of the Bank of Japan's policy of yield curve control, where the yield on Japanese 10-year bonds was pinned at around 0 percent. While Japanese bond yields were already low at the start of our sample period, yields in the UK and the euro area were considerably higher, at around 6 percent in the early 2000s.

Figure 7 shows short-term estimates of expected nominal returns for G4 government bonds. Each panel shows expected nominal returns alongside its two main components: carry and value. The value component dominates the carry component across all markets in our sample period. US government

Figure 5: Impact of bond maturity on estimates of equity risk premium, US data

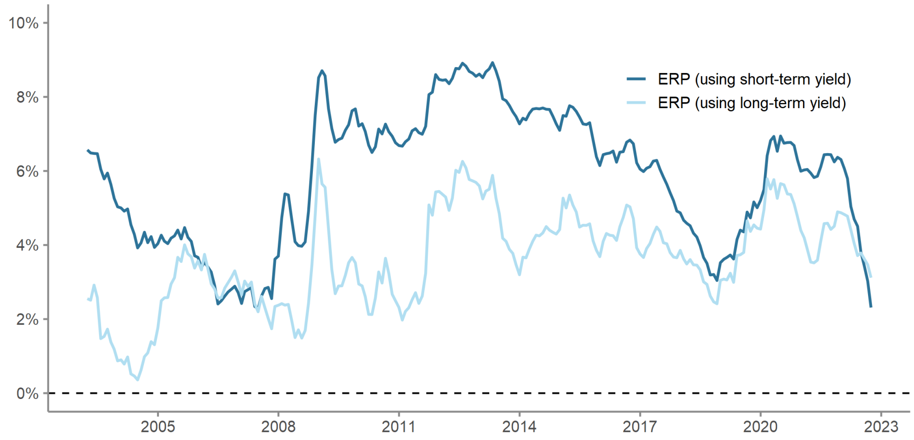

Note: The figure shows two estimates of the expected ERP. Both versions of the ERP are estimated by subtracting expected bond returns from expected equity returns. The estimate labelled "ERP (using short-term yield)" refers to an ERP calculated using 1-year bond yields. The estimate labelled "ERP (using long-term yield)" refers to an ERP calculated using 30-year bond yields. For both estimates, we use expected equity returns based on the framework outlined in this note. The sample period is January 2003 to September 2022.

Figure 6: Long-term expected government bond return estimates

(a) United States *(b) Euro area*

(c) Japan *(d) United Kingdom*

Note: The figure shows estimates of long-term expected government bond returns as defined in equation (11)*. The sample period is January 2003 to September 2022.*

bonds stand out as the market with the most volatile expected returns.

Figure 7: Estimates of expected returns on a 10-year government bond at a one-year horizon

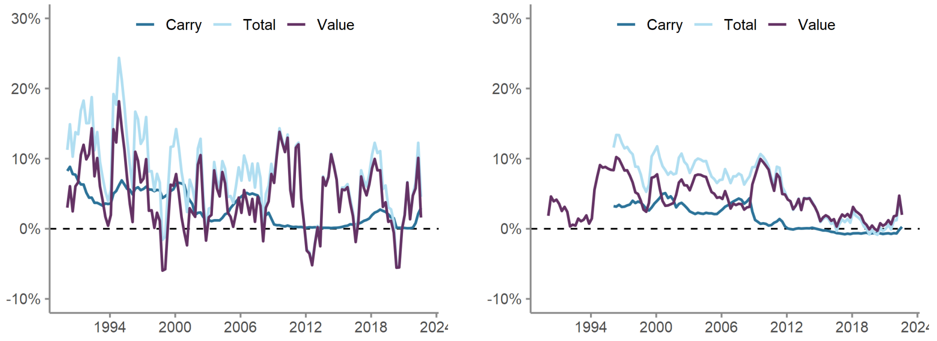

(a) United States *(b) Euro area*

(c) Japan *(d) United Kingdom*

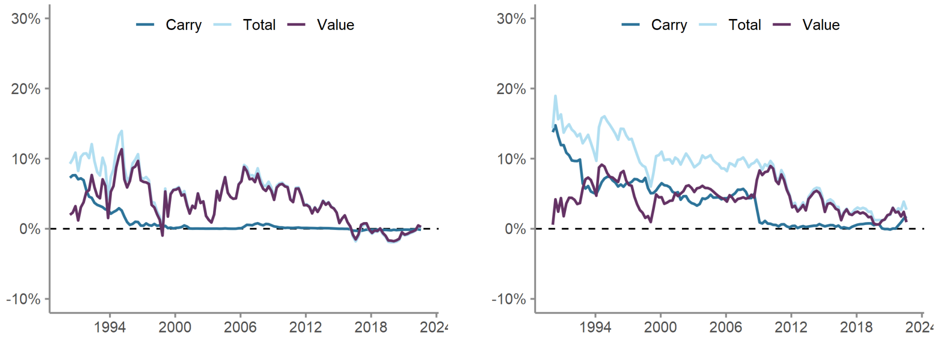

Note: The figure shows estimates of expected returns on a 10-year government bond as defined in equation (12) at a 1-year horizon. The sample period is January 1990 to September 2022. Quarterly frequency.

To obtain aggregate expected government bond returns, we combine expected return estimates for the G4 markets using GDP weights. Figure 8 shows aggregated expected returns on G4 government bonds, including both nominal and real estimates. By the end of Q3 2022, real yields are around 0.7 percent, and both expected nominal and real returns have recovered from close to their sample-lows, which were reached during the outbreak of the Covid-19 pandemic in early 2020. The chart also underscores that negative real yields are not a recent development, as real yields have been negative for almost a decade.

Figure 8: Aggregate long-term expected government bond return estimate

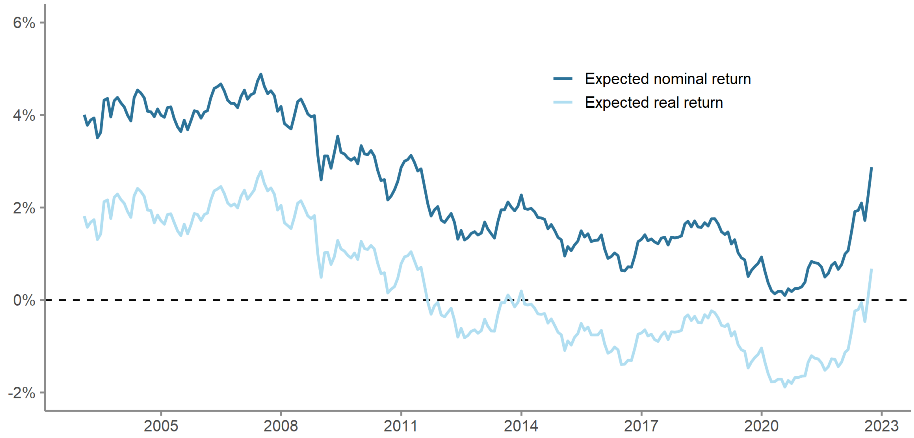

Note: The figure shows GDP-weighted estimates of long-term expected G4 government bond returns as defined in equation (11)*. Expected real returns are obtained by subtracting survey-based long-horizon inflation forecasts from expected nominal returns. The sample period is January 2003 to September 2022.*

Corporate Bond Expected Return Estimates

Figure 9 shows long-term estimates of expected returns on corporate bonds denominated in USD and EUR. Each panel shows expected total returns alongside the expected credit risk premium. Expected total returns on corporate bonds were declining between the GFC and the outbreak of the Covid-19 pandemic. The dominant driver of the recent increase in expected returns has been the increase in government bond yields.

Figure 9: Long-term expected credit return estimate

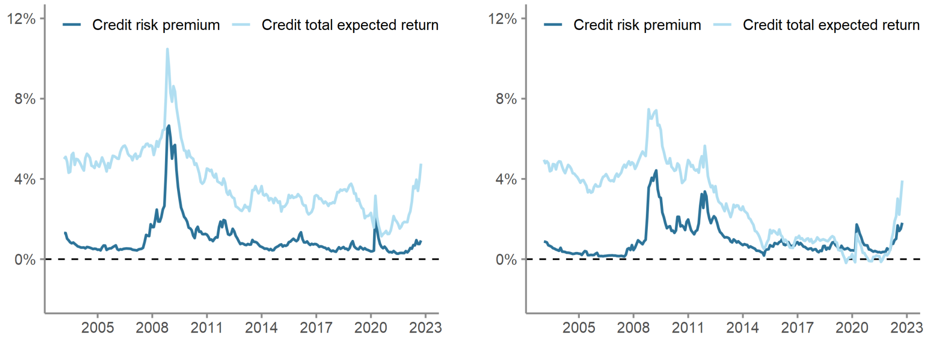
(a) US dollar *(b) Euro*

Note: The figure shows expected corporate bond returns. The sample period is January 2003 to September 2022.

Unlike total expected returns, the credit risk premium does not exhibit any downward trend in our sample period. This suggests that declining expected credit returns are predominantly driven by lower expected government bond returns. Relatively stable estimates of the credit risk premium echo the key takeaway from our estimates of the equity risk premium: falling bond yields have been a key driver of declining expected returns across asset classes.

Figure 10: Aggregate long-term expected credit return estimate

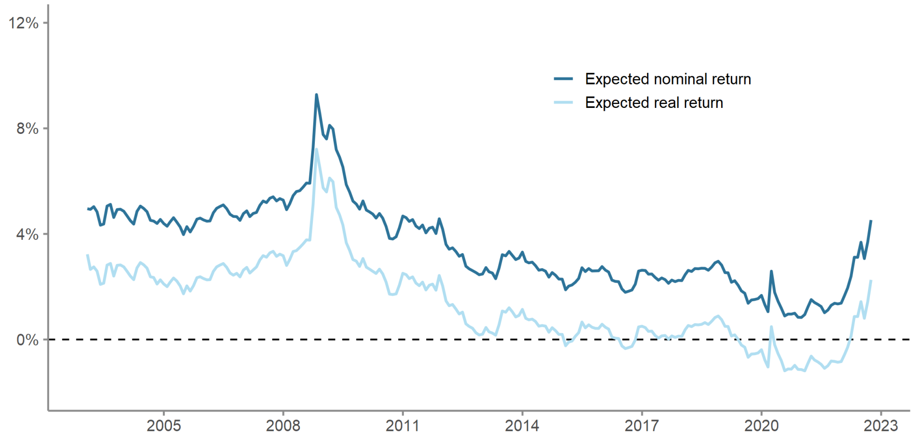

Note: The figure shows market-cap-weighted estimates of expected credit returns. Expected real returns are obtained by subtracting survey-based horizon-matched inflation forecasts from expected nominal returns. The sample period is January 2003 to September 2022.

Figure 10 shows aggregated nominal and real expected corporate bond returns. Similar to expected government bond returns, current estimates of expected nominal returns are close to 4.5 percent, after increasing sharply since the start of 2022. Similarly, expected real returns on corporate credit have seen a sharp increase over the same period, increasing by around 300 basis points to 2.3 percent. As with equities and government bonds, lower yields are the main driver behind the decline in expected corporate bond returns in the pre-pandemic period. Unlike equities, where the impact of falling bond yields has been partially offset by the increasing equity risk premium, estimates of the credit risk premium have remained stable over most of our sample period. This has made the overall decline in expected corporate bond returns since the GFC even larger than those observed for equities and government bonds.

Multi-Asset Expected Return Estimates

We combine expected returns across asset classes to produce expected returns on multi-asset portfolios that approximate the strategic asset allocation of the GPFG. Specifically, we use weights of 70 percent for equities, 21 percent for government bonds, and 9 percent for corporate bonds. The equity allocation is represented by a market-cap-weighted combination of G4 equities. We use a GDP-weighted combination of G4 government bonds for the government bond allocation. The corporate bond allocation is represented by a market-cap-weighted combination of USD- and EUR-denominated corporate bonds.

Figure 11 shows long-term estimates of nominal and real expected returns on the GPFG. The expected return on the fund was declining between the GFC and the outbreak of the Covid-19 pandemic, and has been increasing since then. As at the end of Q3 2022, the expected real return on the fund in local currency is approximately *3 percent*. The corresponding nominal figure is *5.3 percent*.

Figure 11: Estimated expected return on the fund

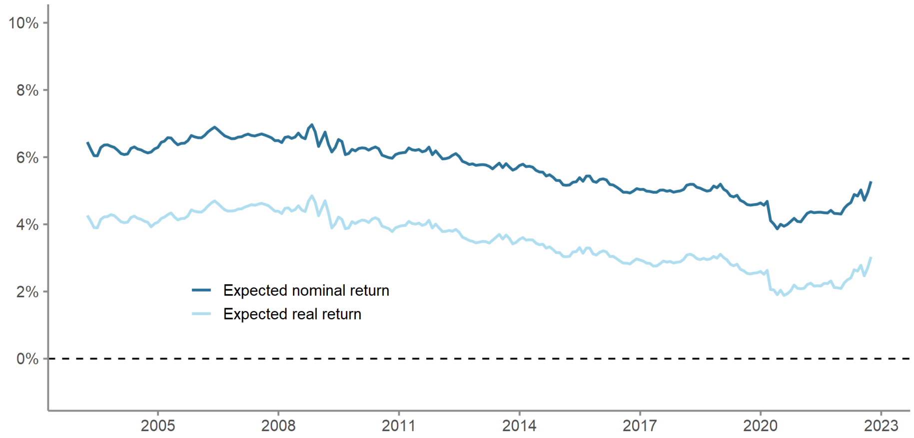

Note: The figure shows estimates of the nominal and real expected return on the GPFG's benchmark in local currency. The sample period is January 2003 to September 2022.

It is worth noting that declining expected returns have coincided with high realised returns across asset classes. While this might appear puzzling at first, it is natural to observe high realised returns on long-duration assets in the presence of shocks to expected returns that are expected to persist for long periods of time – as is the case with the interest rates shocks observed over the last two decades. The relationship arises as declining expected returns (discount rates) lead to higher valuations on long-duration assets such as equities and bonds. As valuations increase, investors realise positive returns on their equity and bond holdings. These valuation gains compensate investors for lower future returns. This dynamic leads to a negative correlation between realised and expected returns on long-duration assets.

5. Forward-Looking vs Backward-Looking Estimates

Our expected return framework is based on the SOP framework, and many approaches to estimating expected returns are essentially versions of this framework – see, for instance, Ferreira and Santa-Clara (2011). Many implementations of this framework rely on historical average returns or predictive regressions, particularly when estimating the growth and value components. The expected return estimates that these models produce are therefore directly or indirectly based on *realised* returns.[20] While all of these implementations aim to capture forward-looking market expectations, several of them end up relying on historical data or assumptions, making them

---

[20] Other implementations use, for instance, survey-based forecasts or more tactical indicators – see Ilmanen (2011) for a comprehensive overview of the most common alternatives across asset classes.

inherently backward-looking.

The framework outlined in this note, by contrast, relies entirely on forward-looking data. To be clear, most of the earlier implementations explicitly caution against extrapolating historical trends into the future. However, several of the forward-looking data sources we draw on in this note were not readily available until a few years ago. Our approach benefits from the availability of these forward-looking data sources, most prominently dividend futures, which contain information about expected dividends at fixed future horizons, and equity index options, which allow us to extract market-implied equity risk premium estimates across horizons. It is therefore useful to compare our forward-looking approach with an alternative implementation that is based on historical returns and fundamentals as estimation inputs.

Figure 12 shows two sets of expected return estimates for US equities. The approach used in this note is labelled "forward-looking" and the alternative implementation using historical data is labelled "backward-looking". The backward-looking version relies entirely on historical returns and past fundamentals as estimation inputs.

The four panels in Figure 12 show estimates of expected equity returns across return components (panels a-c) and in total (panel d). The components labelled as "forward-looking" refer to our estimates. For the backward-looking version, the three components are implemented as follows: "carry" is the 12-month trailing dividend yield, "growth" is the 20-year moving average of a (equally-weighted) combination of earnings and GDP growth, and "value" is the return implied by the price-earnings ratio reverting to its historical mean in a linear fashion over the next ten years. The backward-looking value component is the most often disputed out of the three components, and is sometimes either left out or shrunk towards zero due to its high sensitivity to the chosen sample window. We include it here for completeness and to allow a comparison with our forward-looking value component. These are otherwise standard modelling choices for SOP-style models and should serve as an appropriate alternative to the approach outlined in this note.[21]

As shown in panel a, the two models produce similar estimates of the carry component. This is not surprising as the only potential source of discrepancy is the estimate of next year's index dividend, which can be approximated by the trailing dividend most of the time. There are, however, periods where the two estimates diverge, particularly during market downturns. This divergence can be attributed to the different data sources used to estimate dividends: short-term dividend futures (forward-looking) and 12-month trailing dividends (backward-looking).

The differences between the two approaches become more pronounced for the growth and value components in panels b and c, respectively. Using the backward-looking approach results in estimates of expected cash flow growth that on average are higher than those produced using the forward-looking

---

[21] Rangvid (2017) provides a comprehensive overview of SOP-style models of expected returns.

Figure 12: Forward-looking vs backward-looking expected return estimates, US equities

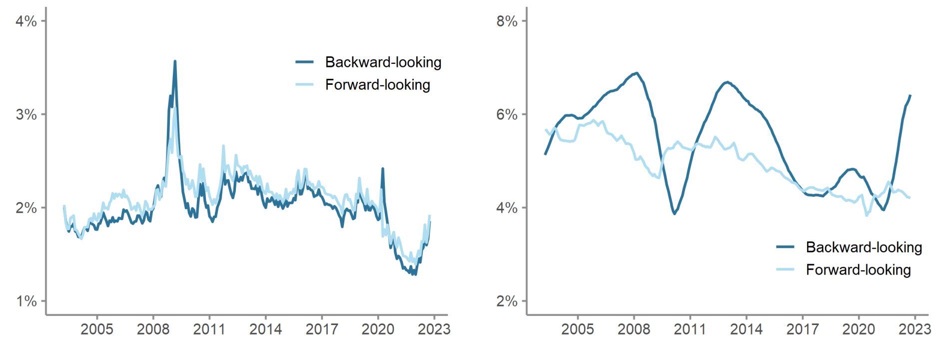

(a) Carry *(b) Growth*

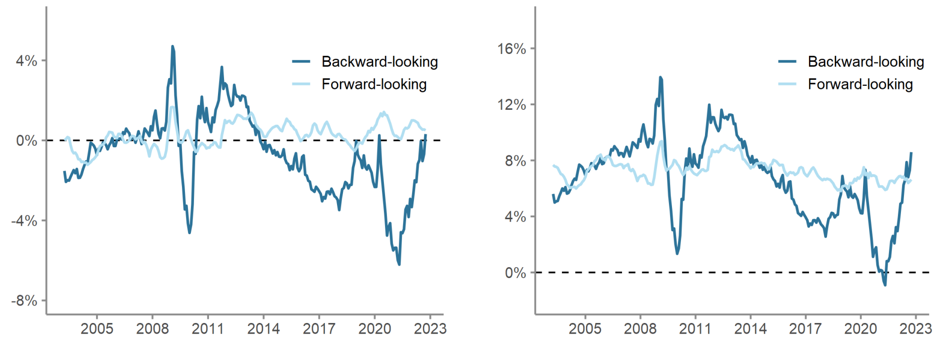

(c) Value *(d) Total*

Note: The figure shows estimates of expected equity returns for each return component (panels a-c) and in total (panel d). The estimates labelled "backward-looking" refer to expected returns based on backward-looking data, where "carry" is the 12-month trailing dividend yield, "growth" is the 20-year moving average of a (equally-weighted) combination of earnings and GDP growth and "value" is the return implied by price-earnings ratios reverting to their historical expanding-window mean (starting in 1980) in a linear fashion over the next ten years. The estimates labelled "forward-looking" refer to expected returns based on the framework outlined in this note. The sample period is January 2003 to September 2022.

version. This difference has averaged at almost 150 basis points over our sample period. In addition, backward-looking estimates of the growth component are significantly more volatile than the forward-looking estimates. The backward-looking estimate suggests that long-horizon growth expectations, which to a large extent determine the growth component, are volatile. If this was the case in the data, the equity present-value model would imply a much higher level of volatility of equity returns.[22]

Finally, the backward-looking estimates track the business cycle with a lag of several years. This is caused by the long moving-average window, which means that backward-looking estimates of expected dividend growth keep trending upwards after a period of high growth, only to slowly decline after growth has peaked. The forward-looking approach, on the other hand, produces cash flow growth expectations that are timely and more stable.

The backward-looking estimates of equity value vary considerably more than

---

[22] Our estimates of expected equity returns, including the growth component, serve as an appropriate benchmark for comparison as they are obtained from a present-value model.

the forward-looking estimates, and are on average negative more often over the sample period. The difference between the two estimates of the value component has been relatively large over the last few years. This happens as the backward-looking estimate compares current equity valuations with the average of past valuations, without taking into account current equity fundamentals, such as the level of interest rates or the equity risk premium. In contrast, forward-looking estimates of the value component are obtained from a present-value model and thus adjust for equity fundamentals, resulting in less dramatic moves in the implied equity value component.

Bringing the three components together in panel d, the estimates of expected equity returns differ both in their absolute level and their variation over time. While the forward-looking estimate at the end of the sample is 6.6 percent, the backward-looking approach gives an estimate of 8.6 percent. Similar to the growth and value components, the estimates of expected equity return derived from the backward-looking model vary considerably more than the forward-looking estimates. A common solution to this problem is to apply shrinkage to reduce the impact from volatile components such as the backward-looking estimates of growth and value. The forward-looking approach arguably offers a more robust alternative to such ad-hoc solutions. We formally evaluate the forecasting performance of both models in Appendix C. Although the sample period is relatively short, the results indicate that the forward-looking model significantly outperforms the backward-looking version.

6. Summary

We estimate expected returns on equities, government bonds, and corporate bonds in developed markets, and combine these to derive expected returns on multi-asset portfolios. We emphasise the use of forward-looking data, predominantly a combination of market-implied expectations and survey-based forecasts, in contrast to traditional approaches based on historical data.

Our estimates indicate that expected returns across asset classes were declining between the GFC (2007-2009) and the outbreak of the Covid-19 pandemic in 2020. As at the end of Q3 2022, long-term expected real returns on developed-market equities and government bonds are estimated to be approximately 3.8 and 0.7 percent, respectively. When combined, this gives an expected real return estimate of 3 percent for a multi-asset portfolio that approximates the benchmark index of the Government Pension Fund Global.

References

Bauer, M. and G. Rudebusch (2020). Interest rates under falling stars. *American Economic Review 110*(5), 1316–1354.

Blanchard, O. (1993). Movements in the equity premium. *Brookings Papers on Economic Activity 2*, 75–138.

Britten-Jones, M., A. Neuberger, and I. Nolte (2011). Improved inference in regression with overlapping observations. *Journal of Business Finance & Accounting 38*(5-6), 657–683.

Cieslak, A. and P. Povala (2015). Expected returns in treasury bonds. *Review of Financial Studies 28*(1), 2859–2901.

Culp, C., Y. Nozawa, and P. Veronesi (2018). Option-based credit spreads. *American Economic Review 108*(2), 454–488.

Desclée, A. and S. Polbennikov (2018). Estimating bond index returns on a long horizon. *Barclays Quantitative Portfolio Strategy Research*.

Elton, E. (1999). Expected return, realized return, and asset pricing tests. *Journal of Finance 54*(4), 1199–1220.

Fama, E. F. and K. R. French (2002). The equity premium. *Journal of Finance 57*(2), 637–659.

Ferreira, M. A. and P. Santa-Clara (2011). Forecasting stock market returns: The sum of the parts is more than the whole. *Journal of Financial Economics 100*(3), 514–537.

Feunou, B. and J.-S. Fontaine (2021). Secular economic changes and bond yields. *Review of Economics and Statistics forthcoming*.

Ibbotson, R. and P. Chen (2003). Long-run stock returns: Participating in the real economy. *Financial Analysts Journal 59*(1), 88–98.

Ilmanen, A. (2011). *Expected Returns: An Investors Guide to Harvesting Market Rewards*. Wiley.

Leibowitz, M. and A. Bova (2012). Duration targeting: A new look at bond portfolios. *Morgan Stanley Research*.

Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics 8*(4), 323–361.

NBIM (2021a). Fundamental drivers of asset returns. *NBIM Discussion Note* (#2-2021).

NBIM (2021b). Modelling equity market term structures. *NBIM Discussion Note* (#3-2021).

Pástor, v. and R. F. Stambaugh (2009). Predictive systems: Living with imperfect predictors. *Journal of Finance 64*(4), 1583–1628.

Rangvid, J. (2017). What rate of return can we expect over the next decade? Working paper.

van Binsbergen, J. (2021). Duration-based stock valuation: Reassessing stock market performance and volatility. Working paper.

Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies 21*(4), 1455–1508.

Appendix A: Expected equity return – derivations

In this appendix, we derive formulas given by equations (6) and (7) in Section 3. We start with the present-value formula for equities:

$$P_t = \sum_{i=1}^{\infty} \frac{E_t D_{t+i}}{\left(1 + R_{t,\infty}\right)^i}.$$

Rearranging, we get:

$$\begin{aligned}
(1 + R_{t,\infty}) &= \frac{E_t D_{t+1}}{P_t} + \frac{E_t D_{t+2}}{P_t \left(1 + R_{t,\infty}\right)} + \frac{E_t D_{t+3}}{P_t \left(1 + R_{t,\infty}\right)^2} + \ldots \\
&= \frac{E_t D_{t+1}}{P_t} + \frac{E_t \left[D_{t+1} \left(1 + \Delta d_{t+2}\right)\right]}{P_t \left(1 + R_{t,\infty}\right)} + \frac{E_t \left[D_{t+2} \left(1 + \Delta d_{t+3}\right)\right]}{P_t \left(1 + R_{t,\infty}\right)^2} + \ldots,
\end{aligned}$$

where $\Delta d_t = D_t / D_{t-1} - 1$. Noting that the price of a dividend strip with maturity n denoted as $P_t^{(n)}$ is given by:

$$P_t^{(n)} = \frac{E_t D_{t+n}}{\left(1 + R_{t,\infty}\right)^n},$$

and defining the weight of each dividend strip in the price of the equity index as:

$$w_t^{(n)} = \frac{P_t^{(n)}}{P_t} = \frac{E_t D_{t+n}}{P_t \left(1 + R_{t,\infty}\right)^n},$$

we can rewrite the equation above as follows:

$$\begin{aligned}
R_{t,\infty} &\approx \frac{E_t D_{t+1}}{P_t} + w_t^{(1)} E_t \Delta d_{t+2} + w_t^{(2)} E_t \Delta d_{t+3} \ldots \\
&\approx \frac{E_t D_{t+1}}{P_t} + \underbrace{\sum_{n=2}^{\infty} w_t^{(n-1)} E_t \Delta d_{t+n}}_{\text{growth annuity } G_{t,\infty}}.
\end{aligned}$$

Appendix B: Corporate bond expected credit loss methodology

This section describes the methodology we use to estimate the expected credit loss on investment-grade corporate bonds. To construct the measure, we need a product of the probability of default and the loss given default.

We follow the methodology outlined in Culp, Nozawa, and Veronesi (2018) to estimate the probability of default. It is based on modelling empirically observable balance sheets of "pseudo firms". Specifically, we assume that these pseudo firms hold publicly traded equity indices as their assets that are funded by a combination of equity and zero-coupon corporate bonds. Using a no-arbitrage condition, risky corporate debt can be replicated with a portfolio of a risk-free debt and a short position in a put option written on the pseudo firm's asset:

$$\hat{B}_t = K\hat{Z}_t\left(T\right) - P_t^{eq.idx}\left(K,T\right), \tag{16}$$

where $\hat{Z}_t$ is the risk-free discount factor at time t with maturity T, and $P_t^{eq.idx}$ is the value of a put option on an equity index at time t with strike price K and maturity T.

Given that we observe both the pricing of risk-free bonds and put options on equity indices, we are able to empirically back out the market value of a pseudo corporate bond and its corresponding metrics such as credit spread or the expected default probability. We use equity index options written on the S&P 500 and EURO STOXX 50 indices to represent corporate bonds denominated in USD and EUR, respectively. We source the option data from OptionMetrics.

On any given date, we have a range of options with different strike prices. Each level of a put strike represents a pseudo firm with a different leverage ratio denoted as K_i / A_t, where A_t is the value of a pseudo firm's assets at time t. This means that we have a range of pseudo corporate bonds with varying expected default probabilities for horizon T which we estimate on a monthly frequency in the following way:

$$\hat{p}_{i,t}(T) = Pr[A_{t+T} < K_i|F_t]. \tag{17}$$

We estimate $\hat{p}_{i,t}(T)$ using the historical return distribution of A_t available at time t. For the USD and EUR pseudo bonds, we use historical S&P 500 and EURO STOXX 50 index cumulative return distributions over horizon T on an expanding-window basis. We compare these returns to a pseudo firm's leverage ratio to empirically estimate how often it would have had its asset values lower than its liabilities (i.e., would have defaulted) at the end of horizon T. We focus on a horizon of $T = 2$ years as equity index option maturities do not extend beyond 2 years.

For investment-grade bonds, we compare estimated default probabilities with historical default rates based on Moody's long-term data. To account for the business cycle variation in default rates, we use different default rates for "expansions" and "recessions". We follow the NBER business cycle dating to identify recessions and expansions. Specifically, if we are in an "expansion" phase, we use the average historical default rate calculated over the "expansion" phase as a basis for comparison. We conduct an analogous procedure for the "recession" period. To construct an investment-grade bond portfolio, we take all pseudo bonds with estimated default probabilities lower than the average historical default rates and aggregate them on an equal-weighted basis.

We pair estimated default probabilities with the historical loss-given-default rate for investment-grade bonds as informed by Moody's corporate default and recovery rates study. Specifically, we use a recovery rate approximated to 45 percent to compute the loss-given-default rate. Similarly to the expected default probabilities, we differentiate the recovery rate along the "expansion" and "recession" dimension. Culp, Nozawa, and Veronesi (2018) estimate that during expansions the recovery rate is 5 percent higher than the unconditional average, whereas during recessions it is 27 per cent lower. As a result, we adjust the unconditional recovery rate in line with these estimates.

Appendix C: Forecasting equity returns

Table 1 compares the forecasting performance of the backward-looking and the forward-looking version of expected US equity returns. If both models are perfectly accurate representations of expected equity returns, the estimated intercept is zero and the beta coefficient on the expected return is one. While both models are imperfect, the estimated coefficients for our forward-looking version are closer to the desired values.

The results indicate that our forward-looking model outperforms the backward-looking version both in and out of sample, as indicated by the adjusted R^2 and by the MSE ratio, respectively. Both versions outperform the naïve benchmark in the form of average realised return, with the forward-looking version having an out-of-sample R^2 that is double that of the backward-looking version.

The results presented in Table 1 should be treated with caution as the sample period is relatively short, especially when considered in relation to the return horizon of 5 years.

Table 1: Forecasting 5-year returns on US equities (2003-2022)

	Backward-looking	Forward-looking
Intercept	0.13*	−0.16
	(2.32)	(−0.83)
Expected return	−0.63	3.32
	(−0.81)	(1.51)
N	189	175
adj. R^2	−0.00	0.15
adj. R^2_{oos}	0.11	0.30
$MSE^m_{oos}/MSE^{bwd}_{oos}$	1.00	0.79

*Note: The sample period is January 2003 to September 2022, monthly data. The dependent variable is annualised total US equity returns. t-statistics are reported in parentheses and are adjusted for overlapping observations following Britten-Jones, Neuberger, and Nolte (2011). * indicates significance at $p < 0.05$. R^2_{oos} is calculated as $1 - MSE^m_{oos}/MSE^{bm}_{oos}$ where MSE^m_{oos} and MSE^{bm}_{oos} are the MSE of the model in question and the benchmark model, respectively (see e.g. Welch and Goyal, 2008). The benchmark model is based on average realised US equity returns computed over an expanding window, starting in 1980. $MSE^m_{oos}/MSE^{bwd}_{oos}$ is the MSE ratio between the model in question and the backward-looking model.*