Norges Bank Investment Management

#1 | 2023

Discussion note

Macro trends and long-horizon returns

We outline a macro-finance model for simulating long-horizon returns on government bonds and multi-asset portfolios. We show how the determinants of return distributions change across different investment horizons, and discuss differences in prospective returns on long- and short-duration bonds.

Date 23/08/2023

dn@nbim.no
www.nbim.no

Summary

- We outline a model for simulating forward-looking return distributions for short and long-duration government bond and multi-asset portfolios, over different investment horizons. The model anchors the yield curve to expectations of long-term inflation and output growth - which we refer to as “macro trends” - and allows for uncertainty around these trends going forward.
- Yields on government bonds have been declining over the last four decades. A large part of the decline in yields can be attributed to declining macro trends, generating high returns on government bonds, especially those with long duration. Looking ahead, the evolution of macro trends will be a key driver of long-term returns. We show how the impact of trends on returns differs depending on bond duration and investment horizon.
- Our framework allows us to make quantitative comparisons of prospective long-term returns. Using a realistic calibration, we show that it is unlikely that long-duration bond returns will match the experience in the last few decades. When assuming that trends are flat on average, the return distributions of short- and long-duration bonds are comparable, despite long-duration bonds earning a term premium. For long-duration bonds to generate returns closer to historical experience, long-term growth prospects would likely need to deteriorate from today’s levels.
- We extend the model to include equity prices alongside the yield curve, and simulate multi-asset return distributions for portfolios with short- and long-duration bonds. When assuming a negative correlation between equity and bond returns, long-duration bonds lower the volatility of a multi-asset portfolio more than short-duration bonds. This improves this distribution of returns over long horizons. These diversification benefits can be large enough to counter the effects of higher macro trends. This result partly depends on the negative equity-bond correlation, however, and the benefit from long-duration bonds is smaller when the correlation is positive or zero.

1. Introduction

Over the past several decades, there has been a secular decline in long-term government bond yields. This can be attributed to declines in long-term inflation and growth expectations, which together we refer to as “macro trends”. Falling yields and macro trends have led to high historical returns on government bonds, with longer-duration bonds outperforming short-duration bonds by a significant margin over this period.

In this note, we simulate prospective long-term returns on short- and long-duration bonds. Since the high realised returns on government bonds have resulted from a secular decline in macro trends, historical returns are a limited guide to the future. We therefore use a forward-looking framework that simulates the distribution of macro trends, where we explicitly model expectations of long-term inflation and output growth. These expectations are used to model the equilibrium interest rate perceived by investors. We model other key drivers of bond yields, such as interest rate cycles and term premiums, and a ‘real rate gap’ - a wedge between the equilibrium real rate and long-term growth expectations. We calibrate our model to US macroeconomic data and asset prices, and generate realistic long- and short-duration fixed income returns. We use the model to simulate asset returns over investment horizons up to 20 years.

We use the simulation model to understand the determinants of fixed income returns over different investment horizons. We show that the evolution of macro trends is the main driver of long-horizon returns, for both short- and long-duration bonds. At longer horizons, trends are a more important driver of returns than the term premium, which is traditionally the main focus when comparing bonds of different maturities. The impact of macro trends is not the same across bonds of different durations, however. Short-duration bonds are more exposed to the level of macro trends, which determines return compounding over long horizons. In contrast, long-duration bonds hedge changes in macro trends, where declines in trends increase the prices of these bonds.

We use our framework to make quantitative comparisons of prospective bond returns given today’s trend levels as a starting point. Given the importance of the path for macro trends for bond returns, we consider alternative calibrations guiding how they evolve *on average* over the long term. In the ‘Baseline’ calibration, macro trends remain near today’s levels, on average, though the distribution of possible outcomes is still very wide. We also consider ‘Normalisation’ and ‘Low Growth’ calibrations where macro trends on average increase or decrease, respectively.

For the Baseline and Normalisation calibrations, the relative performance of long- and short-duration bonds looks different to historical experience. In the Baseline case, long-duration portfolios perform comparably with short-duration portfolios over long horizons, despite earning a positive term premium. In the Normalisation case, long-duration bonds are unlikely to outperform, as returns on short-duration bonds compound higher macro trends. There are higher returns on long-duration portfolios in the Low Growth calibration, which is closest to the experience over the past few decades. Here, we assume that macro trends decline on average, implying that long-term growth prospects deteriorate from today's levels.

We also simulate return distributions for multi-asset portfolios. We compare portfolios with 70% in equities and 30% in fixed income, where we vary the duration of the fixed income component. Long-duration fixed income assets add diversification benefits when the correlation is negative, leading to a lower volatility of multi-asset portfolios and better long-term performance. These diversification benefits can be large enough to counter the effects of normalising macro trends. This result is more finely balanced, however, when assuming the equity-bond correlation is positive or close to zero.

The note proceeds as follows. Section 2 describes the historical drivers of bond yields and returns. Section 3 outlines the simulation framework, describing macroeconomic processes and how they are incorporated into yield curve and equity pricing. In Section 4, we describe the model parameters and initial values, and the different calibrations for the long-term averages of macro trends. In Section 5, we compare distributions of long- and short-duration fixed income returns using the different calibrations. In Section 6, we compare long- and short-duration fixed income when viewed in a multi-asset context. Section 7 concludes.

2. Historical drivers of bond yields and returns

In this section, we provide background on the drivers of the decline in yields over the past several decades. We describe how macro trends have been key drivers of bond yields and how this has translated into higher returns on longer duration bonds. These considerations motivate the design of the simulation model later in the note.

Macro trends in government bond yields

We can use an accounting identity for yields to distinguish between different components influencing their historical decline.[1] The yield on a nominal n-period government bond, $y_t^{(n)}$, can be decomposed according to the following identity:

$$y_t^{(n)} = \underbrace{r_t^* + \pi_t^*}_{i_t^*} + \frac{1}{n}\sum_{j=0}^{n-1} E_t\left(\bar{r}_{t+j}\right) + \frac{1}{n}\sum_{j=1}^{n} E_t\left(\bar{\pi}_{t+j}\right) + tp_t^{(n)},$$

where r_t^* is the long-term or "equilibrium" real rate and π_t^* denotes long-term inflation expectations.[2] The cyclical real rate and inflation components, $\bar{r}_t$ and $\bar{\pi}_t$ respectively, are defined relative to the long-term expectations. In addition to interest rate expectations, yields also embed a maturity-specific term premium, $tp^{(n)}$. Figure 1 shows the 10-year US Treasury yield over the last several decades, alongside measures of the equilibrium real interest rate and long-term inflation expectations.

[1] We previously discussed the drivers of bond yields in detail in NBIM (2021a).

[2] The equilibrium real rate refers to the concept of a natural rate of interest, the rate that brings output into line with its potential level in the absence of transitory shocks or nominal frictions.

FIGURE 1 Decomposition of 10-year nominal US government bond yield

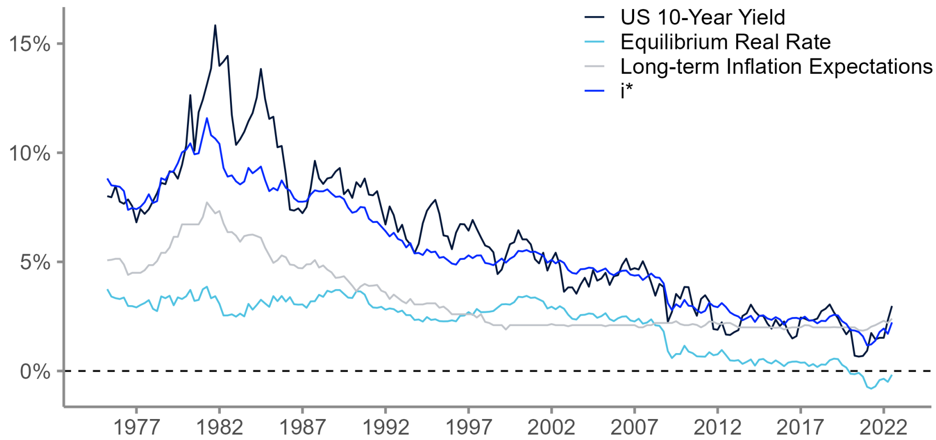

NOTE: The chart shows the nominal 10-year US Treasury yield alongside "i^*", the sum of long-horizon inflation expectations and the equilibrium real rate. The equilibrium real rate splices NY FED estimates (1961-2017) with survey-based estimates (2018-2022) and long-horizon inflation expectations are from the Federal Reserve Board US model. Sample period from Q1 1973 to Q2 2022.

The sum of the equilibrium rate and long-horizon inflation expectations, or i^*, has closely tracked the long-term government bond yield over this period.[3] Broadly speaking, the long-term decline in yields can be attributed to declines in the two long-term components across the two halves of the sample.[4] Earlier in the sample, long-term inflation expectations steadily declined, which is often attributed to monetary policymakers tackling high inflation in this period and stabilising macroeconomic volatility and expectations. Later in the sample, further declines in i^* can be attributed to a falling equilibrium real rate, in particular around the Global Financial Crisis (GFC) in 2007-2009.[5]

Throughout this note, we emphasise the role of long-term inflation and growth expectations - which we refer to as "macro trends" - in accounting for the level of yields over the long term. The equilibrium rate can be linked both theoretically and empirically to long-term growth expectations.[6] Figure 2 shows a measure of r_t^* alongside a measure of long-term growth expectations.[7] The estimate of the equilibrium real rate closely tracks the estimate of potential GDP growth, with the exception of the post-Global Financial Crisis period, which saw a sizeable and persistent gap. We refer to this gap as the "real rate gap", which in our analysis is a wedge between macro trends and the equilibrium short rate, i_t^*. We can therefore

[3] Equilibrium short rate i_t^* is equivalent to the terminal rate that is often discussed in the context of the monetary policy.

[4] For a detailed exploration of the role of trends in term structure models, see Cieslak and Povala (2015) and Bauer and Rudebusch (2020).

[5] Other developed markets experienced declines in government bond yields that were similar in magnitude.

[6] A useful reference is the neoclassical growth model which implies that the natural rate of interest varies over time in response to the output growth rate and shifts in preferences. More details on modelling and estimating the equilibrium real rate can be found in Holston, Laubach, and Williams (2017).

[7] There is a large literature examining the estimation and determinants or the equilibrium interest rate. For an overview, see Kiley (2020). In general, it is challenging to obtain precise estimates of the equilibrium real rate. Our focus in this note is on representing the uncertainty from the investor's perspective rather than obtain the most accurate estimate of the equilibrium real rate.

attribute the decline in bond yields primarily to long-term declines in macro trends and the more recent opening of a real rate gap, and we incorporate these factors into the simulation framework.

FIGURE 2 Equilibrium real interest rate and US potential GDP growth estimate

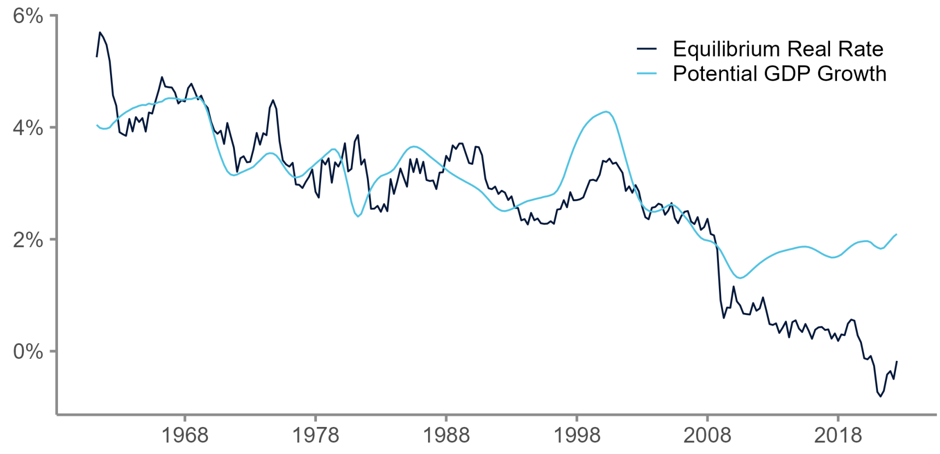

NOTE: "Potential GDP Growth" refers to the growth rate of real potential US GDP from the Congressional Budget Office (CBO). The equilibrium real rate splices NY FED estimates (1961-2017) with survey-based estimates (2018-2022). The survey-based estimate subtracts 10-year expectations of inflation from the 10-year expectations of the three-month Treasury bill rate. The sample period is from Q1:1961 to Q2:2022.

Trends and historical fixed income returns

The steady decline in yields has translated into a sustained period of strong bond returns. Figure 3 compares the cumulative total return from investing one dollar in short- and long-duration US Treasury bonds since January 1976. The cumulative return on long-duration bonds has been considerably higher than the return on a portfolio of short-duration bonds.

A standard approach to comparing short- and long-term bond returns is to consider the slope of the yield curve and term premiums of longer-term bonds. To the extent that the yield curve tends to be upward-sloping, and the term premium positive, we would expect long-term bonds to outperform short-term bonds. However, long-duration bond returns are also more sensitive to changes in macro trends, and have therefore been boosted by the decline in the trends over the last four decades.

These trend effects account for a large proportion of the outperformance relative to short-term bonds over the sample period. Unless we expect to see the same pattern of falling trends again, we should not expect similarly high returns on long-term bonds to repeat. In addition, the opening of the real rate gap will have further increased returns historically, and future returns will depend on whether the gap persists going forward.

Given the limits of using historical returns, we rely on a simulation model for exploring the prospective returns on government bonds. A simulation model

FIGURE 3 Cumulative returns on short- and long-duration US Treasury bonds, 1976-2022

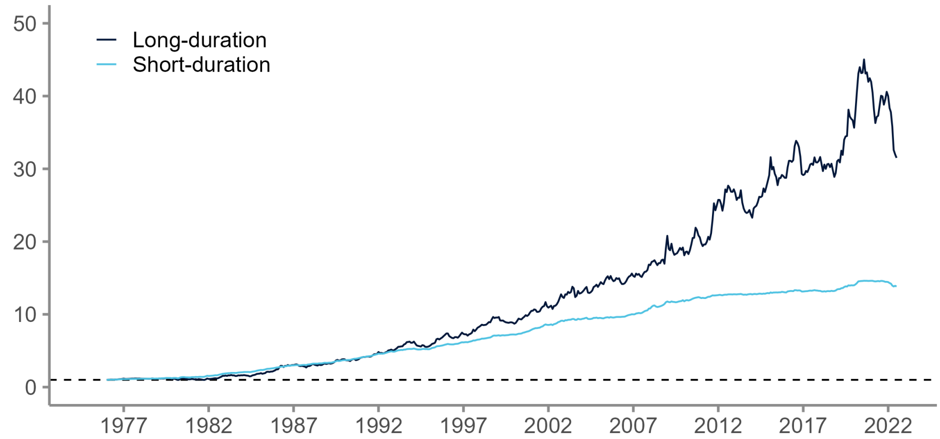

NOTE: Short-duration index includes US Treasury bonds with maturities between one and five years and long-duration index includes Treasury bonds with maturities longer than 20 years. Source: Bloomberg Indices. Sample period from January 1976 to June 2022.

allows us to capture the dynamics of asset prices and macro variables in a forward-looking manner, and to explore a wide range of outcomes for the various return drivers.

3. Modelling trends and asset prices

Next, we describe the model that we use to simulate macro trends, yields, and return distributions for bonds with different durations. We use a reduced-form approach to modelling macro trends and asset prices. At the core of the model are processes for output growth and inflation, that are modelled with persistent and transitory components. These macroeconomic processes provide the foundation for simulating long-term inflation and growth expectations, which evolve slowly over time as investors learn from macroeconomic data. These long-term expectations are then key inputs into modelling the yield curve and equity prices.

Output growth and inflation

The processes for real output growth and inflation are specified in terms of persistent and transitory components. The realised change in the log level of real output, denoted by z_t, is modelled with persistent and transitory components, τ_t^z and a_t, respectively:

$$z_t = \tau_t^z + a_t \tag{1}$$

$$\tau_t^z = (1 - \rho_z)\,\mu_z + \rho_z \tau_{t-1}^z + \varepsilon_t^z \tag{2}$$

$$a_t = \varepsilon_t^a, \tag{3}$$

The persistent component of output growth, τ_t^z, follows an autoregressive process with high persistence captured through a value of ρ_z close to one. The transitory component of growth, a_t, is independent across time with a zero mean, and the long-term average of output growth is given by μ_z.[8]

Inflation, denoted as π_t, is similarly modelled with persistent and transitory components:

$$\pi_t = \tau_t^\pi + c_t \tag{4}$$

$$\tau_t^\pi = (1 - \rho_\pi)\,\mu_\pi + \rho_\pi \tau_{t-1}^\pi + \varepsilon_t^\pi \tag{5}$$

$$c_t = \rho_c c_{t-1} + \varepsilon_t^c. \tag{6}$$

Both the persistent and transitory components of inflation follow autoregressive processes. Persistence is generated by ρ_π close to one, and the parameter ρ_c captures additional persistence in historical inflation.[9] We include additional persistence in the transitory component of inflation, and not for transitory growth to capture the higher persistence of inflation in historical data.

We use an approach to modelling output and inflation that is common in studies that forecast these variables. This is a simple approach that is suitable for simulating over long horizons. The persistent components of output growth and inflation are stationary, which means that both macro variables have constant long-term averages, μ_z and μ_π. These levels are important determinants of asset returns over the long run. We therefore later explore three alternative calibrations of the long-term averages of output and inflation.[10] We describe these calibrations in more detail in the next section.

Expectations of long-term output growth and inflation

Investors do not observe the persistent and transitory components underlying the inflation and growth processes described above. They only observe realised values of z_t and π_t, which they use to infer the long-term dynamics of each macro variable, in particular μ_z and μ_π.

Following an extensive literature modelling macro expectations, we assume that investors apply an adaptive learning rule to infer the properties of the data-generating process from realised values of output and inflation.[11] While the long-term averages μ_z and μ_π are constant, their perceived values vary over time due to learning about the persistent component. A realistic representation of how investors form their expectations about key macro variables is important for modelling asset prices, especially when there is a long duration of cash flows.

We model investors' long-term expectations for each macro variable using a

[8] Various versions of the reduced-form two-component model of real output growth are common in the literature, see e.g., Mueller, Stock, and Watson (2020).

[9] Related studies that represent inflation as a two-factor process include Stock and Watson (2007); Faust and Wright (2013).

[10] An alternative approach would be to make long-term averages time-varying. To keep our model simple, we represent the variation in the long-term averages through three alternative calibrations. By exploring these, it is easier to convey the key intuitions.

[11] See, for example, Evans, Honkapohja, and Williams (2010). The adaptive learning approach has been motivated in terms of learning from lifetime experiences, for example applied to inflation in Malmendier and Nagel (2016) and asset prices in Nagel and Xu (2022). Earlier literature motivated adaptive expectations in terms of their optimality in the presence of regimes and other sources of instability that render the data from distant past less relevant.

constant-gain learning rule.[12,13] Investors apply the learning rule to obtain real-time estimates of the persistent components as follows:

$$\bar{\tau}_t^z = \nu_z \bar{\tau}_{t-1}^z + (1 - \nu_z)\, z_t \tag{7}$$
$$\bar{\tau}_t^\pi = \nu_\pi \bar{\tau}_{t-1}^\pi + (1 - \nu_\pi)\, \pi_t. \tag{8}$$

The long-term expectations, $\bar{\tau}_t^z$ and $\bar{\tau}_t^\pi$, update in response to realised macro data, z_t and π_t. The learning parameters, ν_z and ν_π, determine the speed of updating. These parameters are close to one, which means that investors update their estimates of $\bar{\tau}_{t-1}^z$ and $\bar{\tau}_{t-1}^\pi$ slowly. The long-term expectations in our model correspond to the macro trends described in Section 2. In our model, the estimate $\bar{\tau}_t^\pi$ can be directly interpreted as a measure of π_t^*.[14]

Based on the macro trends, investors also perceive transitory components of macro variables that differ from the underlying components, a_t and c_t. We define these real-time estimates of transitory components as:

$$\bar{a}_t = z_t - \bar{\tau}_t^z \tag{9}$$
$$\bar{c}_t = \pi_t - \bar{\tau}_t^\pi. \tag{10}$$

An important consequence of investor-learning is that there can be a non-trivial wedge between the long-term expectations and the actual unobserved persistent components. These wedges can persist over long horizons, where the parameters of the underlying macro model are difficult to learn even with decades worth of data.[15]

Equilibrium real interest rate

The process for the equilibrium real interest rate, r_t^*, combines expectations of output growth with two additional variables.[16] As shown in Figure 2, the equilibrium real rate co-moves with long-term growth expectations, and this is a feature of a wide range of macroeconomic models that incorporate a standard intertemporal consumption problem. There are deviations between the two series at different points in time, however, in particular following the 2008/09 financial crisis. In the model, expectations of long-term growth are directly represented by $\bar{\tau}_t^z$. In

[12] We define long-horizon expectations as follows: $\bar{\tau}_t^z \equiv \lim_{i\to\infty} E_t\left[\tau_{t+i}^z\right]$ and $\bar{\tau}_t^\pi \equiv \lim_{i\to\infty} E_t\left[\tau_{t+i}^\pi\right]$.

[13] Constant-gain learning rules share similarities with full-memory Bayesian learning models, see Farmer, Nakamura, and Steinsson (2022) for a recent reference. Unlike these models, however, it ensures that the learning effects never disappear.

[14] If a central bank targets inflation, $\bar{\tau}_t^\pi$ can be interpreted as the perceived inflation target (Kozicki and Tinsley, 2001).

[15] The fact that the trend components are unobservable has several important implications for the behaviour of asset prices and how they relate to macroeconomic variables, see e.g. Collin-Dufresne, Johannes, and Lochstoer (2016); Nagel and Xu (2022) for a more detailed discussion.

[16] Similar to the learned components of macro variables, we refer to the equilibrium rate that is perceived by investors, as opposed to any 'true' equilibrium rate for the economy.

addition, we include two processes that allow for deviations between r_t^* and $\bar{\tau}_t^z$:

$$r_t^* = \bar{\tau}_t^z - s_t - \gamma_t \tag{11}$$

$$s_t = (1-\rho_s)\mu_s + \rho_s s_{t-1} + \epsilon_t^s \tag{12}$$

$$\gamma_t = (1-\rho_\gamma)\mu_\gamma + \rho_\gamma \gamma_{t-1} + \epsilon_t^\gamma \tag{13}$$

The variable s_t represents a "safety" component of interest rates, that captures the documented convenience yield in US Treasury bonds.[17] We include an additional variable, γ_t, which we refer to as the "real rate gap". This variable captures a wider set of factors that could account for the low level of the equilibrium rate, beyond a low level of long-term growth.[18] While we have an incomplete understanding of the real rate gap, it can potentially be an important driver of yields and other asset prices. In light of this, later in the note we consider alternative scenarios for the persistence of the gap, and their implications for our results.

Monetary policy and the yield curve

Our modelling of the yield curve builds on the macro trends and equilibrium real rate. We specify a no-arbitrage term structure model that includes macro trends, where the short-term interest rate is determined by a monetary policy rule.[19] The policy rule characterises the behaviour of the short-term nominal interest rate, i_t, as follows:

$$i_t = i_t^* + \phi_c \bar{c}_t + \phi_a \bar{a}_t + \phi_i \bar{i}_{t-1} + u_t^i, \tag{14}$$

where $i_t^* = \pi_t^* + r_t^*$ and $\bar{i}_t = i_t - i_t^*$. The rule draws from a large literature on modelling monetary policy, where simple rules have been shown to be able to accurately describe monetary policy rates over time. The short-term rate is set in line with transitory inflation and output growth, where the coefficients ϕ_c and ϕ_a describe the relative weights placed on output and inflation stabilisation.[20,21] We include persistence in the monetary policy rule through the loading ϕ_i on the lagged interest rate cycle, $\bar{i}_{t-1}$, and also include a monetary policy shock, u_t^i.

The process for the short-term interest rate incorporates macro trends and

[17] This component captures the narrowly defined safety component that is usually identified as a spread between the government bond and a bond that does not have the safety/liquidity features of a government bond but is otherwise identical to it. Such a safety component is a permanent feature of US Treasury bonds (Krishnamurthy and Vissing-Jorgensen, 2012). The reason for including the safety component in the equilibrium real rate as opposed to including it as part of the term premium is the fact that the safety component has a roughly equal impact across all maturities whereas the relative importance of term premium increases with the maturity.

[18] Recent studies such as Davis, Fuenzalida, and Taylor (2021) and Brand, Goy, and Lemke (2021) take a similar approach in modelling r^* and growth. There are many possible explanations for this gap, with key candidates being asset purchases from quantitative easing (Koijen, Koulischer, Nguyen, and Yogo, 2021), or a re-assessment of risk following the financial crisis that increased demand for risk-free assets (Kozlowski, Veldkamp, and Venkateswaran, 2020).

[19] Recent related work in this area include Bauer and Rudebusch (2020), Favero, Melone, and Tamoni (2021), and Feunou and Fontaine (2021).

[20] We specify the monetary policy rule in terms of output growth, rather than in terms of the output gap level. This modelling choice is natural based on our modelling of macro variables, and there is theoretical and empirical support in various studies, see for example Walsh (2003), Orphanides (2003), Coibion and Gorodnichenko (2011) and Coibion and Gorodnichenko (2012).

[21] We assume that investors and the central bank share the same information set, where the central bank also learns about long-term growth and inflation and must set policy on this basis.

cyclical components, from which we build the yield curve. In addition to these variables, yields include a term premium which we represent through a one-factor structure:

$$x_t = (1 - \rho_x)\,\mu_x + \rho_x x_{t-1} + \varepsilon_t^x. \tag{15}$$

The average level of term premium is determined by μ_x. We collect the cyclical components of bond returns in the vector $\bar{X}_t = (\bar{i}_t, \bar{c}_t, \bar{a}_t, x_t)^{'}$. The n-period nominal zero-coupon yield, $y_t^{(n)}$, is given by:

$$y_t^{(n)} = i_t^* + a_n + b_n^{'}\bar{X}_t, \tag{16}$$

where a_n and b_n are recursions determined by the no-arbitrage condition imposed in the term structure model. We calibrate the factor loadings in the model to match historical yields, which we show in Appendix A.

Modelling equity prices

In addition to simulating the yield curve, we model equity prices using macro trends and yield curves as inputs. Following the methodology outlined in NBIM (2021b), we model the value of the aggregate stock market index, S_t, and the corresponding index dividend, D_t. The price of the index dividend paid out n years from now, denoted by $P_t^{(n)}$, is the present value of D_{t+n}:

$$P_t^{(n)} = D_t \exp\left(n\left(g_t^{(n)} - y_t^{(n)} - \theta_t^{(n)}\right)\right), \tag{17}$$

where $g_t^{(n)}$ is the annualized expected dividend growth at the n-period horizon, $y_t^{(n)}$ is the n-period nominal yield and $\theta_t^{(n)}$ is the risk premium compensating investors for dividend risk at the n-period maturity. The value of the equity index is the sum of the present values of all future dividends:

$$S_t = \sum_{n=1}^{\infty} P_t^{(n)}. \tag{18}$$

The index price is therefore a function of three term structures - expected dividend growth, the yield curve, and risk premiums. We model dividend growth in line with nominal output growth in the model:

$$g_t = \tau_t^z + \tau_t^\pi + (a_t + c_t)\,, \tag{19}$$

where $g_t^{(n)}$ is the expected value of g_t n periods ahead. We model risk premiums, θ_t, also in terms of persistent and transitory components:

$$\theta_t = \theta_t^* + \tilde{\theta}_t \tag{20}$$

$$\theta_t^* = \mu_{\theta^*} + \beta_{\theta^*} x_t \tag{21}$$

$$\tilde{\theta}_t = \rho_{\tilde{\theta}}\tilde{\theta}_{t-1} + \varepsilon_t^{\tilde{\theta}}, \tag{22}$$

where θ_t^* and $\tilde{\theta}_t$ are the long-term and cyclical risk premiums, respectively. The three term structures allow us to obtain the price of equities as a present value of expected dividends discounted with the risk-adjusted discount rate.[22]

Following NBIM (2021b), we implement the present value model for equities as a two-stage model. In the first stage, which covers the first 30 years, we model individual dividend strips given by equation (17). The remaining equity value is represented through a perpetuity. The total equity return is the sum of the equity price change and the dividend paid out over the period considered.

4. Model calibration and alternative trend paths

Next, we describe how the model is calibrated to match macroeconomic and asset price data, and how we set initial values for the simulations. This exercise involves setting the model parameters and defining a covariance matrix of shocks, which we calibrate using US data at a quarterly frequency.

We consider three alternative calibrations that differ in terms of the long-term averages of output growth and inflation, μ_z and μ_π. These parameters are key determinants of the long-horizon distributions of macro variables and asset prices. The average values of macro trends will be determined by these parameters, but there is considerable uncertainty around what these values are. To capture this uncertainty, our alternative calibrations allow for high and low long-term averages of real GDP growth and inflation, shown in Table 1.

TABLE 1 Long-term average values (%) of macro variables, historical vs. alternative calibrations

Variable	Baseline	Normalisation	Low Growth	Historical
Real GDP Growth (μ_z)	1.8	2.9	0.2	2.9
Inflation (μ_π)	2.3	3.2	1.1	3.2

NOTE: Historical averages for Normalisation are estimated using quarterly data from Q1 1960 to Q3 2022. Data are sourced from FRED.

For the 'Baseline' calibration of the model, we set the long-term averages of macro trends in line with Consensus Economics long-term forecasts of output growth and inflation. These values are similar to the current estimates of macro trends, meaning that on average macro trends are broadly unchanged in the Baseline calibration. We also consider a 'Normalisation' calibration, where the long-term average levels of macro trends are set in line with historical averages of output and inflation. These higher values imply that macro trends increase on average over the simulation horizons. Finally, we consider a 'Low Growth' alternative where macro trends fall on average over the simulation horizon.[23]

For each alternative calibration, we set the other model parameters in order to match key historical moments. Table 2 shows the historical and simulated standard

[22] For simplicity, we assume that equities are real assets, which means that inflation in the discount rate and on the cash flow side cancels out. This assumption can be easily relaxed by changing the loading on inflation in equation (19), to better align with the empirical evidence.

[23] We set the unconditional averages of the Low Growth calibration to match survey forecasts of long-term growth and inflation for Japan as of Q4 2022 from Consensus Economics.

deviations and autocorrelations of output growth, inflation, macro trends and asset returns. We report additional details on the model calibration in Appendix A.

TABLE 2 Moments of macro variables and asset returns, historical vs. simulated (annualised)

Variable	σ (%, Hist.)	σ (%, Simulated)	AC (Hist.)	AC (Simulated)
Panel A. Macro Variables				
Real GDP Growth	2.4	2.2	0.78	0.77
Inflation	2.3	2.2	0.97	0.93
Panel B. Macro Trends				
Growth Trend ($\bar{\tau}_t^z$)	0.8	0.8	0.99	0.99
Inflation Trend ($\bar{\tau}_t^\pi$)	1.9	1.6	0.99	0.99
Panel C. Bond Yields				
1Y Bond	3.5	2.3	0.96	0.97
10Y Bond	3.0	2.4	0.97	0.94
Panel D. Asset Returns				
1Y Bond	2.4	1.4	0.39	0.68
10Y Bond	12.5	16.4	-0.05	0.08
30Y Bond	36.9	32.9	-0.05	0.00
Equities	15.8	15.7	-0.02	-0.06

NOTE: σ and AC refer to standard deviation and first-order autocorrelation, respectively. Observed moments are estimated using quarterly data over the period from Q1 1967 to Q3 2022. Sample period starts in Q1 1962 for one- and 10-year yield and in Q2 1977 for the remaining yields. Sample period ends in Q3 2022. Data are sourced from FRED.

Table 2 shows we are able to simulate volatilities and persistence of macro variables and trends in line with historical moments. As reported in Appendix A, we also generate a positive correlation between output growth and inflation. We use the calibrated macro variables as inputs into the calibration of the yield curve and equities. The model captures the properties of asset returns, generating moments of returns in line with realised estimates for bonds of different maturities and equities. We set the long-term average equity premium equal to 4%, and generate a negative correlation between the 10-year government bond returns and equity returns of -0.30.[24] Additional details on the parameters for the yield curve and equities are provided in Appendix A, where we also outline the calibration of monetary policy, convenience yield, and risk premiums.

In each of the calibrations, we assume that the real rate gap closes steadily over several years (the long-term average, μ_γ, equals zero). We calibrate a half-life for the gap of approximately 7 years, and by setting the variance of the gap process at a low level, we ensure that the gap consistently follows this path across different simulations. To capture uncertainty around the path for the real rate gap, we later consider a scenario where the gap remains constant over the investment horizons we consider in our analysis.

[24]We explore the implications of a positive correlation in Section 6.

Simulation initial values

We set the initial values for the simulation model based on macroeconomic data and the yield curve in Q4 2022. We aim to match the observed 1-, 5- and 10-year nominal US Treasury yields at the end of Q4 2022 as closely as possible. For macro trends, we set the values of $\bar{\tau}_t^z$ and $\bar{\tau}_t^\pi$ in line with the Baseline calibration of long-term averages, at 1.8% and 2.3%, respectively. We set the short-term interest rate i_t equal to 4.50%, and the initial value of the term premium equal to 0.70%.[25]

The value of $\bar{\tau}_t^z$ feeds into the starting value for the equilibrium real rate, r_t^*. For the convenience yield, s_t, we use an initial value of 0.50%, in line with its long-run average. We set the value of the real rate gap, γ_t, equal to 100 basis points, which gives an initial value of r_t^* equal to 0.3%. The perceived transitory macro variables, $\bar{a}_t$ and $\bar{c}_t$, are set at -0.25% and 1.75%, respectively. At the end of Q4 2022, the yields on 1-, 5- and 10-year US Treasuries were 4.7%, 4.0% and 3.9%, respectively. With the set of starting values we have described, the initial levels of the 1-, 5- and 10-year yields in the simulation model are 4.6%, 4.3% and 3.8%, respectively.

5. Distributions of fixed income returns

In this section, we use the simulation model to analyse the distributions of fixed income returns over different horizons. From the initial values described above, we simulate 10,000 alternative paths for macroeconomic variables and yield curves, up to horizons of 20 years at a quarterly frequency.

Using the simulations, we first discuss the distributions of macro trends over different horizons. We show how different yield components determine the distribution of returns on short- and long-duration bonds over different horizons. We then compare simulated paths for returns on short- and long-duration bonds under the alternative calibrations of macro trends.

Macro trends and return drivers over different horizons

Given the importance of macro trends highlighted so far, we start by showing the distribution of combined expectations of long-term growth and inflation across horizons. Figure 4 shows the simulated distributions at horizons ranging from one to 20 years, using the Baseline model calibration.

At shorter horizons, the range of outcomes is relatively concentrated around the initial starting points for the macro trends. Over long horizons, under the Baseline calibration, the long-term averages of the macro trends are in line with their initial values, but the distributions are very wide. This dispersion highlights the risk of persistent changes in macro trends from their current level. The wide range of outcomes for macro trends covers a broad set of macroeconomic regimes. Even under the Baseline calibration, the distributions are wide enough to assign a significant probability of low growth and low inflation environments as well as

[25] The short term rate is in line with the top of the target range for the Federal Funds Rate in Q4 2022. The value for the term premium is set in line with internal model estimates.

FIGURE 4 Distribution of macro trends across investment horizons

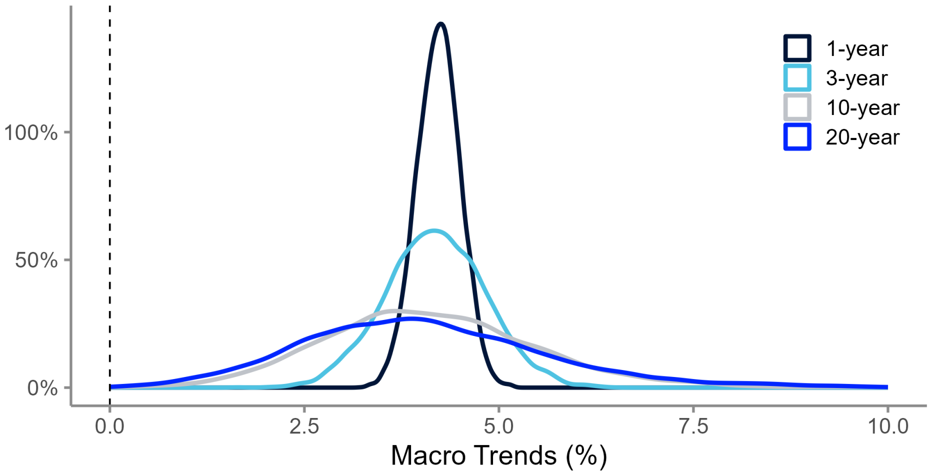

NOTE: The chart shows the distribution of the sum of long-term inflation expectations, $\bar{\tau}_t^z$, and long-term growth expectations, $\bar{\tau}_t^\pi$, across horizons generated by the Baseline calibration of the model.

those with high inflation.[26] This implies that there are significant long-term risks within the model, which are important drivers of asset prices. The distributions of macro trends across horizons illustrate the nature of long-term risk investors face.

Given our focus on distributions of fixed income returns, a key question is what role macro trends play in determining returns. We are interested in how this role changes by return horizon, and how it compares to other yield curve components. Using the model, we can quantitatively show how macro trends influence fixed income returns depending on bond duration and investment horizon.

The average return path is determined by the level of bond yields, which in turn depend on the level of macro trends, term premiums, the real rate gap, and other components of yields. Each of these components varies over time, and can generate return paths that are different to the average path. To understand what accounts for this distribution of returns, we decompose the variance of returns into contributions from different components. Figure 5 shows contribution of return drivers across horizons up to 20 years. Panels (a) and(b) decompose 1-year and 10-year bond portfolios, respectively.[27]

We divide return drivers into 'Rate cycle', 'Term premium' and 'Macro trends' categories. We exclude the real rate gap from the decompositions, as the variance across simulations is very low.[28] The chart decomposes the variance of returns across simulations over horizons from 1 to 20 years. Given that we are concerned with multi-period returns, for each return driver we identify two effects: 'Compounding' and 'Repricing'. 'Compounding' captures the variation in returns

[26] For example, the distribution of macro trends presented in Figure 4 includes stagflationary environments with persistently high inflation and low growth.

[27] These refer to short- and long-duration portfolios that maintain a constant duration of one and ten years by rebalancing the portfolio to the target duration every quarter.

[28] The liquidity component of yields could also be included in the set of return drivers, but we omit this from the decompositions. We choose to do this as the contribution to return variance is very small for both maturity bonds, and the liquidity component is not a important factor when comparing returns of long- and short duration portfolios later in the note.

FIGURE 5 Contributions to 1-year and 10-year bond return variance at different horizons

(A) 1-year bond **(B)** 10-year bond

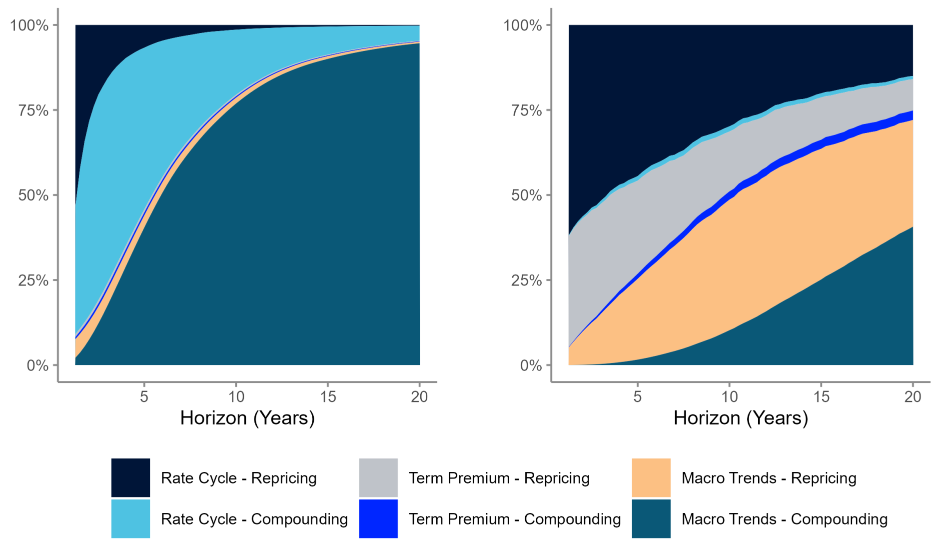

NOTE: Areas show proportion of return variance attributed to drivers. Covariance terms in the variance decomposition are omitted, and proportions rescaled to sum to one.

due to different rates of return compounding. For example, 'Macro Trends - Compounding' describes how returns over long horizons vary depending on whether we continually re-invest in bonds over periods where the levels of macro trends are high or low. 'Repricing' refers to variability in returns that arises from changes in the prices of bonds, for example price increases due to lower macro trends or term premiums.

Figure 5 highlights differences between short- and long-duration fixed income returns. At shorter horizons, for the 1-year bond, the rate cycle components account for almost all return variation. For the 10-year bond, changes in the term premium and rate cycle variation account for most of the variation in returns over short horizons. As the horizon increases, the decompositions change significantly. Since macro trends evolve slowly, they initially play a minor role in explaining variation in both the 1-year and 10-year yields at short horizons. Their importance grows as the horizon increases, however. For the 1-year bond, over the long-term, returns are primarily driven by the compounding of macro trend levels. The short-term strategy effectively re-invests frequently, and so lower (higher) levels of macro trends will lead to lower (higher) return compounding. For the 10-year bond, the contribution of macro trend compounding also increases with horizon, but an equally important driver is repricing due to macro trends. Given the higher duration of long-term bonds, persistent changes in macro trends lead to significant capital gains/losses. The compounding and repricing effects of macro trends are negatively correlated, which captures the ability of long-term bonds to hedge 're-investment risk' from changing interest rates.[29]

[29] This implies there is a negative covariance term required for Figure 5 Panel (b) to sum to the total vari-

For long-term bonds, changes in term premiums decline in importance over longer horizons, similar to the contribution of cyclical variation in yields. The compounding of the term premium increases at longer horizons, but this is not a large contributor given that the persistence of the term premium is lower compared to macro trends. The decompositions highlight the limited role for short-term interest rate cycles in accounting for variation in long-term returns.[30]

Short vs. long duration fixed income returns

Next, we show the distributions of returns of short- and long-duration portfolios over horizons up to 20 years. When interpreting the distributions, we discuss the average path for returns, and the distribution around this path. We use the 'compounding' and 'repricing' concepts described above to understand the *variance* of returns under the Baseline, Normalisation, and Low growth paths for macro trends.

Throughout this section, we show distributions of nominal returns, though our framework also allows us to compare the properties of real returns. When expressed in real terms, our comparisons of returns on short- and long-term bonds across calibrations do not materially change. For completeness, we include distributions of real returns across the different calibrations in Appendix B. We start by exploring the Baseline calibration of macro trends. Figure 6 shows the distributions of cumulative returns for the short- and long-duration portfolios, in Panels (a) and (b), respectively.

FIGURE 6 Cumulative return fans for 1- and 10-year nominal government bonds, Baseline calibration

(A) 1-year bond

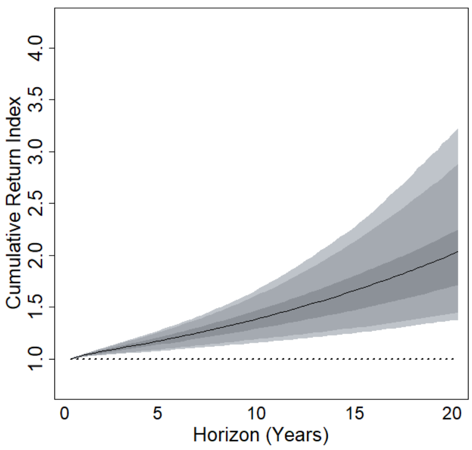

(B) 10-year bond

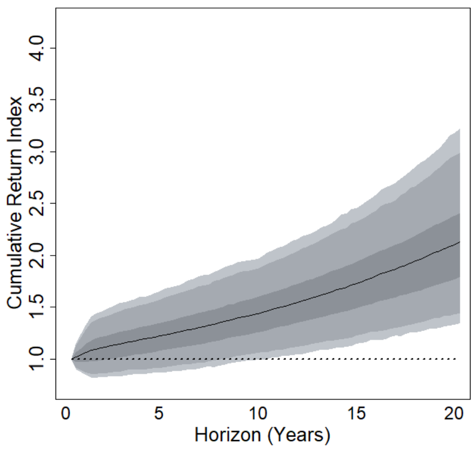

NOTE: Dark band covers 50 percent probability, medium band 90 percent, and light band 95 percent. Black line shows the average value each quarter.

The average path for returns is relatively similar across the two panels. Despite the fact that the long-duration strategy earns the term premium on average, it does

ance of 10-year bond returns at each horizon.

[30] This suggests that when short-term interest rates are constrained near to zero, the effects will be small when considering return variation over long horizons.

not materially outperform the short-duration strategy.[31] This is partly due to the real rate gap steadily closing over the simulation horizon. As the gap closes, expected returns compound at a higher level, for both short- and long-term bonds. While the term premium is positive, it is partly offset by the closure of the real rate gap. Figure 7 compares the distributions of returns at the 3- and 20-year return horizons. Table 3 in Appendix C reports the percentiles of the cumulative distributions at different horizons.

FIGURE 7 Distribution of cumulative returns on 1- and 10-year nominal government bonds at 3- and 20-year horizons, Baseline calibration

(A) 3-year horizon **(B)** 20-year horizon

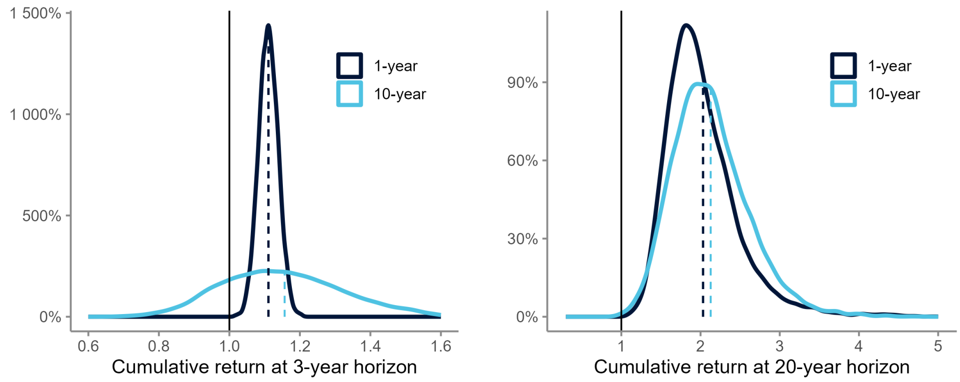

NOTE: Dashed lines indicate mean values of distributions.

At short horizons, the width of the return distributions varies significantly between the short- and long-duration strategies. In Panel (a), the volatility of the shorter-duration strategy is lower, with a lower mean value compared to the long-duration portfolio. The higher long-duration mean at the 3-year horizon results from the real rate gap not closing enough to offset the term premium on average. Long-term bond returns over short horizons are more volatile in general, due to the longer duration of cash flows.

At longer horizons, the return distributions of the two strategies are more comparable, as shown in Figure 7 Panel (b). At the 20-year horizon are similar, the widths of the long- and short-duration distributions are similar. For both bond strategies, there is a wide distribution of 20-year returns, where a significant driver of this is the variability of macro trends. As discussed above, the 1-year bond strategy is exposed to re-investment risk as macro trends change, and this implies that the variance of returns grows with the investment horizon. The 10-year bond strategy hedges this risk to an extent, where declining macro trends imply a positive repricing effect, and vice-versa. This implies that the distribution of returns does not widen by as much for the longer-term bond at longer horizons.

Next, we turn to the Low Growth and Normalisation calibrations, and again compare the distributions of returns for the two bond portfolios. As discussed above, the long-term averages of macro trends play an important role in

[31] We repeat our analysis in Appendix D assuming that the term premium is zero on average. The long-duration strategy would slightly underperform on average in this scenario.

determining long-horizon returns on fixed income. For the Low Growth calibration, macro trends decline on average over the simulation horizon. Figure 8 compares the short- and long-duration strategies under this calibration. Figure 9 shows the return distributions at fixed horizons.

FIGURE 8 Cumulative return fans for 1- and 10-year nominal government bonds, Low Growth calibration.

(A) 1-year bond **(B)** 10-year bond

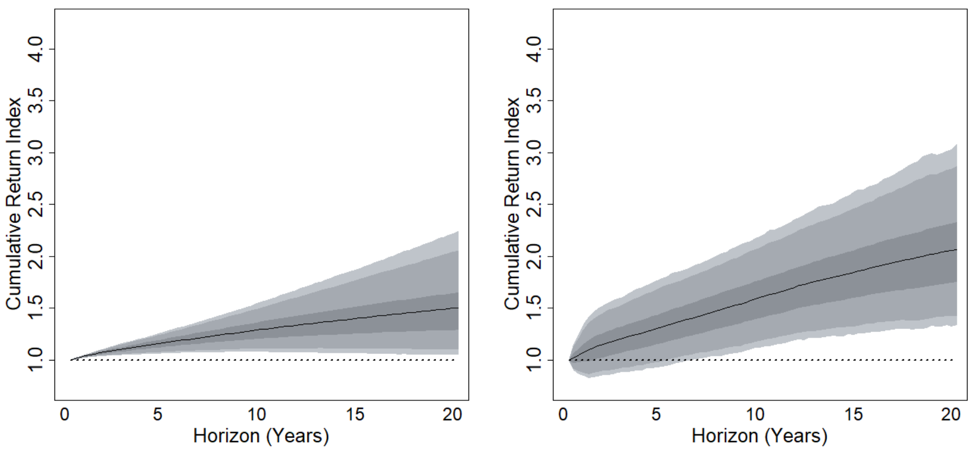

NOTE: Dark band covers 50 percent probability, medium band 90 percent, and light band 95 percent. Black line shows the expected value each quarter.

FIGURE 9 Distribution of cumulative returns on 1- and 10-year nominal government bonds at 3- and 20-year horizons, Low Growth calibration

(A) 3-year horizon **(B)** 20-year horizon

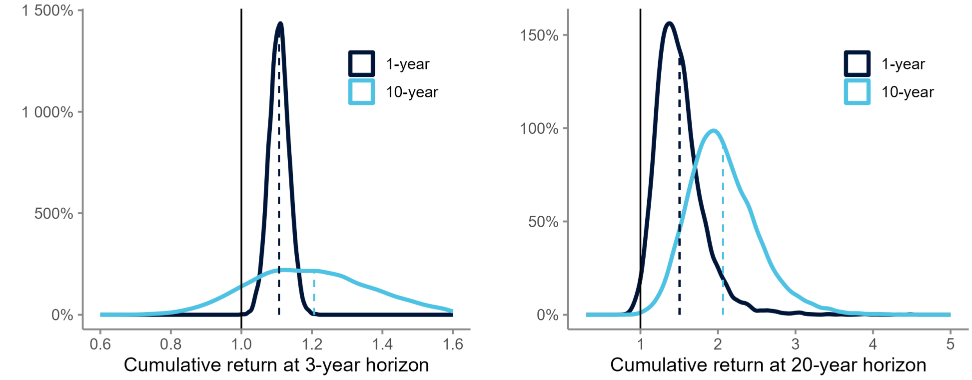

NOTE: Dashed lines indicate mean values of distributions.

Under the Low Growth calibration, at the 3- and 20-year horizons, the mean return is higher for the long-duration portfolio. The distribution of returns on the long-duration portfolio compares favourably due to the gradual declines in macro trends from their current levels.[32] The long-horizon returns of the short-duration portfolio are primarily determined by the compounding of macro trends. The

[32] Table 4 in Appendix C reports the percentiles of the cumulative returns for both alternatives.

lower level of trends means that returns compound at a lower rate, and this means the 1-year bond return distribution moves to the left under the Low Growth calibration. The macro trend declines are associated with positive 'repricing' returns for long-term bonds, which offsets the lower compounding effects. At the 20-year horizon, a significant portion of the long-duration distribution lies to the right of the short-duration return distribution. These results more closely resemble historical experience and the outperformance of long-term bonds shown in Section 2. Under this calibration, the real rate gap again closes slowly over several years. The importance of this effect is diminished, however, where the drag on long-duration returns is more than offset by declining macro trends.

We next turn to the Normalisation calibration, where macro trends increase on average from their current level to a level in line with historical data. Figure 10 compares the short- and long-duration strategy under this calibration, and Figure 11 shows the return distributions at fixed horizons.

FIGURE 10 Cumulative return fans for 1- and 10-year nominal government bonds, Normalisation calibration.

(A) 1-year bond **(B)** 10-year bond

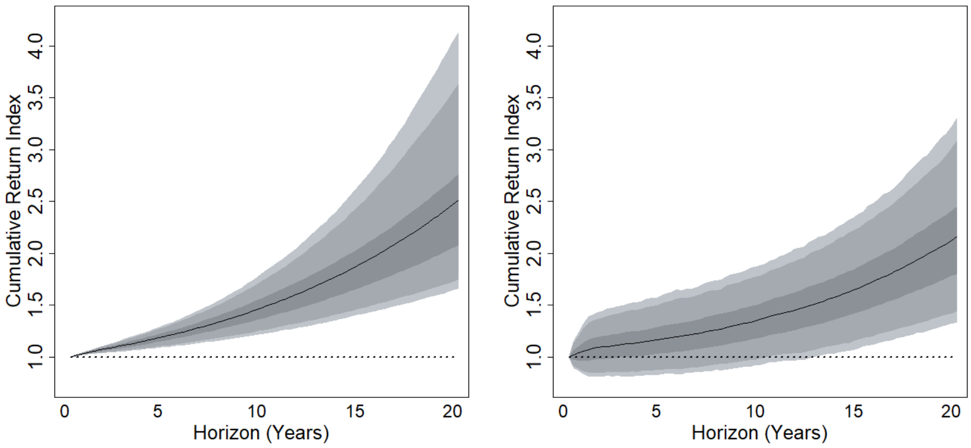

NOTE: Dark band covers 50 percent probability, medium band 90 percent, and light band 95 percent. Black line shows the expected value each quarter.

Compared to the Baseline and Low Growth calibrations, the long-duration portfolio on average underperforms the short-duration portfolio. In the Normalisation calibration, long-duration bonds are more likely to underperform short-duration bonds. The effects have the opposite sign to what we saw in the Low Growth calibration. As macro trends increase, short-duration bonds benefit from re-investing at higher levels of expected returns, and therefore compound at higher rates over long horizons. For long-term bonds, these increasing trends lead to negative 'repricing' returns, which offset the higher rates of compounding. Short-duration bonds are less exposed to changing macro trends, and therefore have less of a drag from higher macro trends in the Normalisation calibration. As a result, the long-duration distribution lies to the left of the short-duration distribution at the 20-year horizon.

FIGURE 11 Distribution of cumulative returns on 1- and 10-year nominal government bonds at 3- and 20-year horizons, Normalisation calibration

(A) 3-year horizon **(B)** 20-year horizon

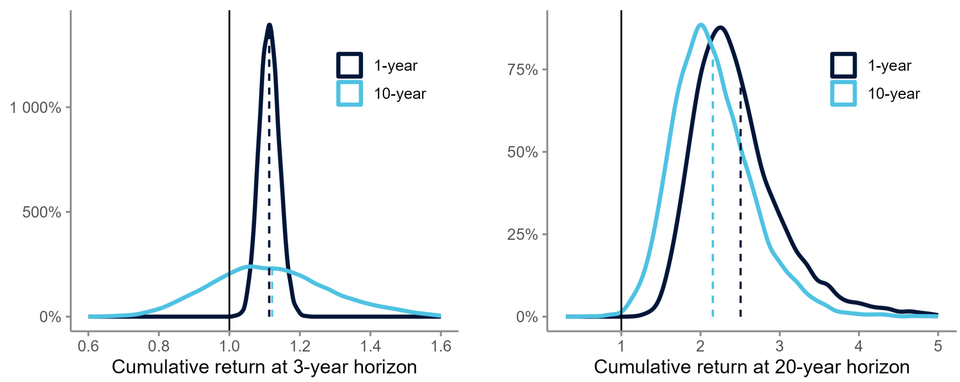

NOTE: Dashed lines indicate mean values of distributions.

To sum up, for the Baseline and Normalisation calibrations, the relative performance of long- and short-duration bonds looks different to historical experience. In the Baseline case, long-duration portfolios perform comparably with short-duration portfolios over long horizons, despite earning a positive term premium. In the Normalisation case, long-duration bonds are unlikely to outperform, as returns on short-duration bonds compound higher yields due to higher macro trends. There are higher returns on long-duration portfolios in the Low Growth calibration, more closely resembling historical experience.

Alternative scenarios for the real rate gap

Under the alternative calibrations so far, we have assumed that the real rate gap closes slowly over several years, which contributes to a higher equilibrium real interest rate. The initial value of the real rate gap implies that the equilibrium real rate is 100 basis points lower than the long-term growth outlook would suggest.

As discussed above, this creates a drag on fixed income returns, in particular for longer-duration bonds.[33] To assess the importance of this assumption, we consider an alternative scenario where the gap is constant over the simulation horizon. We explore the implications of a constant gap for the Baseline calibration only. The conclusions for the Normalisation and Low Growth calibrations do not change when assuming a constant gap.[34]

Figure 12 shows the distribution of cumulative returns on short- and long-duration fixed income portfolios under the constant gap scenario.

[33] Given our limited knowledge about the drivers of the real rate gap, we opt for exploring a range of plausible scenarios for the future evolution of the gap rather than pinning our analysis on one particular calibration.

[34] There may be interesting interactions between the calibrations we have defined and the scenarios for the real rate gap. For example, the increase in the gap around 2008/9 could plausibly be linked to central bank quantitative easing policies, that aimed to stimulate economic growth. In the Low Growth calibration, we might expect central banks to use their balance sheets to lower bond yields over long periods. This could be represented through a more persistent gap in the Low Growth calibration, which we leave for future work.

FIGURE 12 Distribution of cumulative returns on 1- and 10-year nominal government bonds at 3- and 20-year horizons, Baseline calibration with constant real rate gap

(A) 3-year horizon **(B)** 20-year horizon

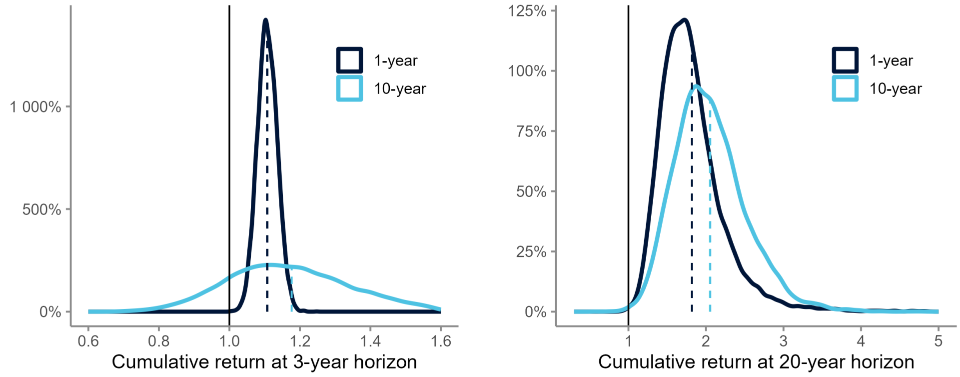

NOTE: Dashed lines indicate mean values of distributions.

The distributions under this scenario differ a small amount from the Baseline results presented earlier. Over the 3-year horizon in Panel (a), the long-duration portfolio no longer experiences a return drag from the closing gap. When the real rate gap persists over the 20-year horizon, in Panel (b), the long-duration portfolio compares more favourably relative to the short-duration portfolio. This is in part due to the short-duration portfolio compounding at a lower rate over long horizons, since yields are lower when the real rate gap persists in this scenario. In addition, by removing the drag on returns for the long-duration portfolio, the term premium is no longer offset over the long term. While the assumption about the path of the gap can alter the relative performance of the fixed income portfolios, the magnitudes of the changes are small compared to the assumptions regarding the path for macro trends.

6. Short vs. long-duration in multi-asset portfolios

Our comparisons of short- and long-duration portfolios have focused on fixed income-only portfolios. Next, we produce distributions of returns for multi-asset portfolios, that combine government bonds with different durations with equities. In standard multi-asset portfolio analyses, the focus is often on the diversification benefits of combining equities and bonds, where the benefits depend on the correlation between the two assets. In our simulations, we can assess how changing the duration of the fixed income portfolio impacts these diversification effects.

We compare two 70-30% equity-fixed income portfolios, where we change the duration of the fixed income portfolio between 1- and 10-years.[35] Given the focus of this note on fixed income returns, we do not decompose the drivers of long-term multi-asset portfolio returns. There are complex interactions between yield components and equity pricing that we plan to examine in future work. In this

[35]We assume continuous rebalancing of the weights of equity and fixed income each quarter.

section, we show the multi-asset return distributions from the model and discuss the effects of changes to the bond-equity correlation. We initially calibrate a negative correlation between equities and bonds, in line with the experience of the last two decades. We then explore the implications of a positive correlation. We use the Baseline calibration of macro trend long-term averages for this analysis.

Figure 13 compares the distribution of cumulative returns for the two alternatives up to a 20-year horizon. The percentiles of the cumulative multi-asset return distributions are reported in Table 6 in Appendix C.

FIGURE 13 Distribution of cumulative returns on multi-asset portfolios at 3- and 20-year horizons, Baseline calibration.

(A) Equity and 1-year bond **(B)** Equity and 10-year bond

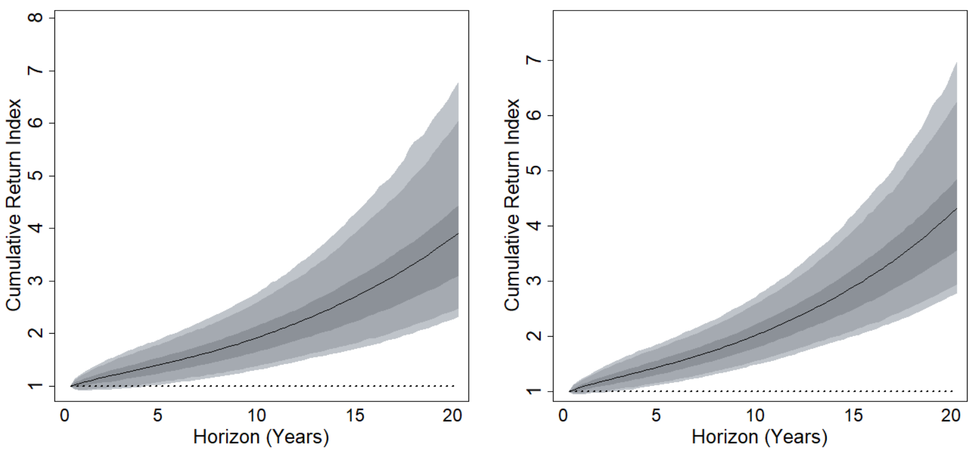

NOTE: Dark band covers 50 percent probability, medium band 90 percent, and light band 95 percent. Black line shows the expected value each quarter.

On average, the cumulative returns on the longer-duration portfolio are higher than the short-duration portfolio. Earlier, using the Baseline calibration, we saw that the long-duration portfolio did not significantly outperform the short-duration portfolio on average over the 20-year horizon. In the multi-asset context, however, the long-duration fixed income portfolio provides greater diversification benefits. This can be seen in Figure 14 Panel (a), which shows the distribution of multi-asset portfolio returns at the 3-year horizon.

At the 3-year horizon, despite the higher volatility of long-term bonds, the multi-asset distribution is narrower compared to portfolio with short-term bonds. This is due to the diversification benefits of the negative correlation between equities and fixed income. The lower volatility of the multi-asset portfolio with long-duration is driven by the fact that long-duration bonds are a more efficient diversifier of equity risk compared to short-duration bonds. Long-duration bonds have a higher volatility, leading to greater hedging benefits when the equity-bond correlation is negative.

Over the 20-year horizon, shown in Figure 14 Panel (b), these diversification benefits lead to an improvement in the distribution of returns on the long-duration

FIGURE 14 Distribution of cumulative returns on multi-asset portfolios at 3- and 20-year horizons, Negative Equity-Bond correlation

(A) 3-year horizon

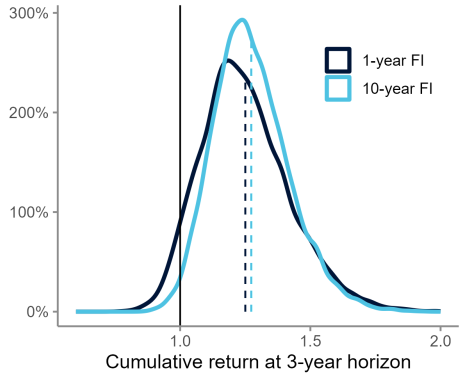

(B) 20-year horizon

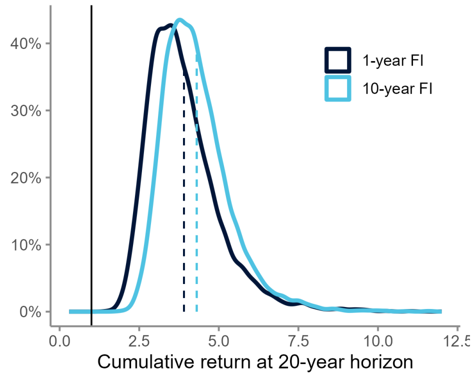

multi-asset portfolio. This implies that, given a negative equity-bond correlation, the case for longer-duration bonds is stronger when viewed from a multi-asset context.

These results also apply for the Low Growth and Normalisation calibrations. Figure 15 shows the 20-year distributions of multi-asset portfolio returns for the Low Growth and Normalisation calibrations, in Panels (a) and (b), respectively. Similar to the fixed income results, lower compounding of expected returns under the Low Growth calibration means that the return distributions shift to the left relative to the Baseline calibration. We also see the opposite effect under the Normalisation calibration. For both calibrations, the return distribution for the long-duration portfolio sits to the right of the short-duration portfolio. For the Normalisation case, the improvement is relatively mild, but the multi-asset diversification effects are able outweigh the adverse effects of trend normalisation that we saw earlier.

FIGURE 15 Distribution of cumulative returns on multi-asset portfolios at 20-year horizon for different calibrations

(A) 20-year horizon - Low Growth calibration

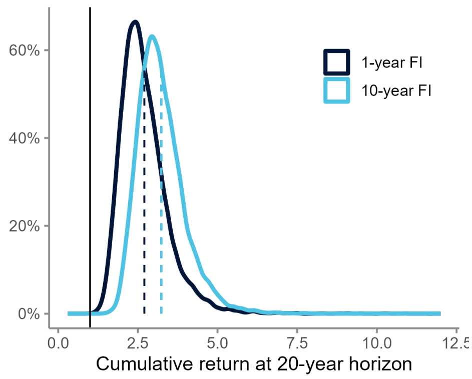

(B) 20-year horizon - Normalisation calibration

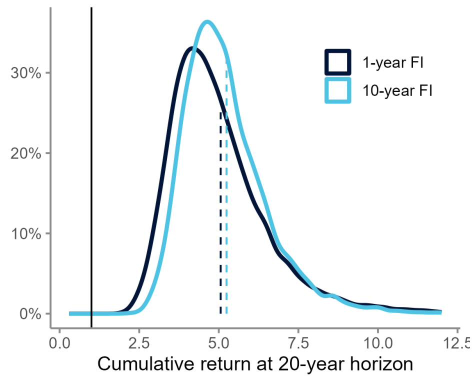

A positive bond-equity return correlation

The sign of the equity-bond correlation is a key parameter affecting multi-asset return distribution. Next, we consider an alternative calibration where the correlation is positive. While such a scenario is likely to be associated with significant changes to many parameters of the simulation model, we make a minimal number of changes in order to switch the sign of the correlation. This allows us to focus on the mechanical effects of a switch in the correlation, rather than pursuing a full scenario analysis.

To generate a positive equity-bond correlation, we change the correlation between real output growth and inflation from positive to negative.[36] This modification changes the inflation behaviour from pro-cyclical to counter-cyclical. This implies that inflation tends to be high when output growth is weak, characterising a stagflationary environment. In addition, we change the correlation between the equity risk premium and the term premium from negative to positive, to reflect the change in risk profile of bonds. The magnitudes of both changes are in line with the empirical evidence from historical inflationary environments in the US.

Figure 16 shows the distributions of cumulative returns for multi-asset portfolios when the equity-bond correlation is positive. Table 7 in Appendix C reports the percentiles of these distributions.

FIGURE 16 Distribution of cumulative returns on multi-asset portfolios at 3- and 20-year horizons, Positive Equity-Bond correlation

(A) 3-year horizon **(B)** 20-year horizon

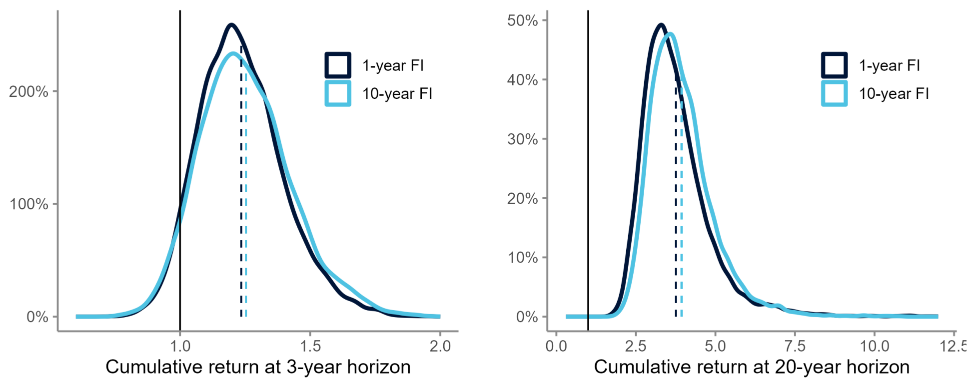

Compared to the negative correlation results, there is less of a difference between the alternative portfolios. At the 3-year horizon, the volatility of returns is now slightly lower when including short-term bonds in the multi-asset portfolio. This is because the higher volatility of long-term bonds is no longer offset by a negative equity-bond correlation. Over the 20-year horizon, there is a smaller difference between the short- and long-duration alternatives.[37] The difference between the short- and long-duration portfolios naturally scales with the value of the

[36] We keep the long-term averages of macro variables unchanged.

[37] As noted earlier, the variance of 1-year returns increases at longer horizons, so the long-horizon distribution will not necessarily sit to the right of the long-duration multi asset portfolio.

equity-bond correlation. In unreported results, we calibrate the model to produce a correlation between equities and fixed income of around zero. Again, there is less of a benefit from long-duration bonds compared to the negative correlation case, though some improvement relative to the case with a positive correlation.

7. Summary

The long-term decline in yields, and corresponding high realised returns on government bonds, can be explained by declining macro trends. These declines have also meant that long-duration bonds have had significantly higher returns compared to short-term bonds over the last few decades. Given that the magnitudes of declines in yields are unlikely to repeat, historical fixed income returns are unlikely to be a good guide to the future. To address this issue when comparing short- and long-term bond returns, we use a simulation model that captures a wide range of possible future trend paths for yields.

Despite long-term bonds earning a positive term premium, their long-term distribution of returns is comparable to short-term bonds. The prospective returns on long-term bonds are therefore likely to undershoot historical performance. Our analysis shows that long-duration bonds are more likely to outperform short-duration bonds when long-term growth prospects deteriorate from today's levels. Longer-duration bonds may provide additional portfolio diversification benefits when viewed from a multi-asset perspective. However, these benefits partly depend on a negative equity-bond correlation persisting over the investment horizon.

References

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

Binsbergen, J., W. Diamond, and M. Grotteria (2020). Risk-free interest rates. *Journal of Financial Economics 143*(1), 1–29.

Brand, C., G. Goy, and W. Lemke (2021). Natural rate chimera and bond pricing reality. ECB Working paper.

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

Coibion, O. and Y. Gorodnichenko (2011). Monetary policy, trend inflation, and the great moderation: An alternative interpretation. *American Economic Review 101*(1), 341–70.

Coibion, O. and Y. Gorodnichenko (2012). Why are target interest rate changes so persistent? *American Economic Journal: Macroeconomics 4*(4), 126–162.

Collin-Dufresne, P., M. Johannes, and L. Lochstoer (2016). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review 106*(3), 664–698.

Davis, J., C. Fuenzalida, and A. Taylor (2021). The natural rate puzzle: Global macro trends and the market-implied r*. NBER working paper.

Evans, G., S. Honkapohja, and N. Williams (2010). Generalized stochastic gradient learning. *International Economic Review 51*(1), 237–262.

Farmer, L., E. Nakamura, and J. Steinsson (2022). Learning about the long run. Working paper.

Faust, J. and J. Wright (2013). Forecasting inflation. In G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 2 of *Part A*, Chapter 1, pp. 3–56. Elsevier.

Favero, C., A. Melone, and A. Tamoni (2021). Monetary policy and bond prices with drifting equilibrium rates. Working paper.

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

Holston, K., T. Laubach, and J. Williams (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics 108*(1), 59–75.

Kiley, M. T. (2020). The global equilibrium real interest rate: Concepts, estimates, and challenges. *Annual Review of Financial Economics 12*, 305–326.

Koijen, R., F. Koulischer, B. Nguyen, and M. Yogo (2021). Inspecting the mechanism of quantitative easing in the euro area. *Journal of Financial Economics 140*(1), 1–20.

Kozicki, S. and P. Tinsley (2001). Shifting endpoints in the term structure of interest rates. *Journal of Monetary Economics 47*, 613–652.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2020). The tail that wags the economy:beliefs and persistent stagnation. *Journal of Political Economy 128*(8), 2839–2879.

Krishnamurthy, A. and A. Vissing-Jorgensen (2012). The aggregate demand for treasury debt. *Journal of Political Economy 120*(2), 233–267.

Malmendier, U. and S. Nagel (2016). Learning from inflation experiences. *Quarterly Journal of Economics 131*, 53–87.

Mueller, U., J. Stock, and M. Watson (2020). An econometric model of international growth dynamics for long-horizon forecasting. *Review of Economics and Statistics*, 1–47.

Nagel, S. and Z. Xu (2022). Asset pricing with fading memory. *Review of Financial Studies 35*(5), 2190–2245.

NBIM (2021a). Fundamental Drivers of Asset Returns. *NBIM Discussion Note* (#2-2021).

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

Orphanides, A. (2003). Historical monetary policy analysis and the Taylor rule. *Journal of Monetary Economics 50*(5), 983–1022.

Stock, J. and M. Watson (2007). Why has u.s. inflation become harder to forecast? *Journal of Money, Credit and Banking 39*(1), 3–33.

Walsh, C. (2003). Speed limit policies: the output gap and optimal monetary policy. *American Economic Review 93*(1), 265–278.

Appendix A: Simulation and calibration details

We collect all processes specified in Section 3 in vector Y_t:

$$Y_t = \begin{bmatrix} \tau_t^z \\ a_t \\ \tau_t^\pi \\ c_t \\ s_t \\ \gamma_t \\ u_t^i \\ x_t \\ \tilde{\theta}_t \end{bmatrix}, \tag{23}$$

which has the following dynamics:

$$Y_t = (I - \Phi)\,\mu + \Phi Y_{t-1} + \epsilon_t. \tag{24}$$

The vector of shocks ϵ_t is normally distributed with:

$$\epsilon_t \sim N\left(0, \Omega_t\right). \tag{25}$$

The covariance matrix Ω_t is a function of lagged state vector Y_{t-1}:

$$\Omega_{t+1} = \Sigma_0 \Sigma_0' + \Sigma_1 Y_t Y_t' \Sigma_1'. \tag{26}$$

This specification of volatility allows for volatility to vary across different macroeconomic regimes. For example, empirical evidence suggests that inflation became more volatile in the high-inflation period of 1970s and 1980s in the US. The specification includes constant volatility as a special case, in which case the model becomes a regular VAR. We generate co-movement across macro variables and economies using the correlation of shocks.

Appendix A.1 Macro variables and equilibrium rate calibration

As discussed in Section 4, we calibrate the parameters of macro processes to match the moments historical real GDP growth and inflation, based on data from Q1 1967 to Q4 2022.

We calibrate the correlation between realised inflation and output growth to be positive, i.e. the inflation is pro-cyclical. This is in line with the empirical relationship from recent decades. In Section 6, we switch the correlation between realised inflation and output growth to negative to explore a stagflationary regime.

The realised values for output and inflation are also used as inputs into the learning rules for $\bar{\tau}_t^z$ and $\bar{\tau}_t^\pi$. The constant-gain learning rule in equations (7)–(8) can be approximated through a discounted moving average with a fixed look-back period. We calibrate learning parameters ν_z and ν_π and the corresponding look-back periods to maximize the correlation with long-horizon surveys. We use the perceived inflation target rate (PTR) from the Fed for inflation and a long-horizon forecast of real output growth from the Consensus Economics panel.

For the convenience yield, we calibrate the s_t process to match the moments of the estimated convenience yield following the methodology outlined in Binsbergen, Diamond, and Grotteria (2020). We set the average value in line with the historical average of 40 basis points. For the real rate gap, we set μ_γ equal to zero and a half-life of approximately 7 years. We set the variance of the ϵ_t^γ equal to near-zero, implying that the real rate gap follows the same path across all simulation paths.

Appendix A.2 Yield Curve and Equities

For the monetary policy rule, we calibrate parameters in line with estimates for the model specified in equation (14) over the period 1983 to 2022. While it is well-documented that these coefficients can change over different monetary policy regimes, the exact parameter values for the monetary policy rule are relatively less important for our results, especially for long horizon returns.

For the term premium process, we calibrate the parameters based on the estimate from Cieslak and Povala (2015). We set the unconditional mean of the term premium equal to 65 basis points. The historical fit of the yield curve model is shown below:

The calibrated macroeconomic processes and yield curves are used as inputs into the calibration of equities in the model. The processes for the risk premium components are calibrated to match the empirical estimates in NBIM (2021b). We use the parameters of the risk premium processes to calibrate the average, volatility and autocorrelations of equity returns. In addition, we target a negative correlation between equity and fixed income returns.

FIGURE 17 Yield curve fit, US data

(A) 1-year

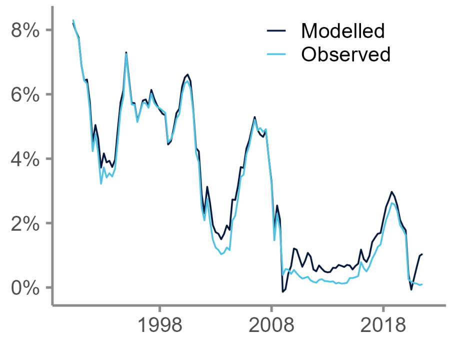

(B) 5-year

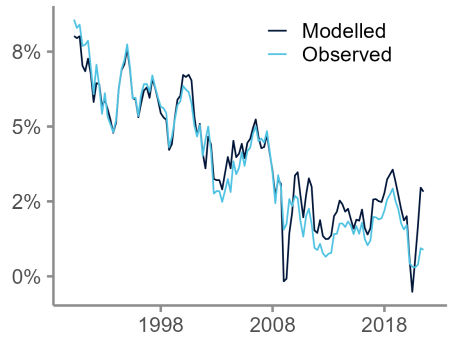

(C) 10-year

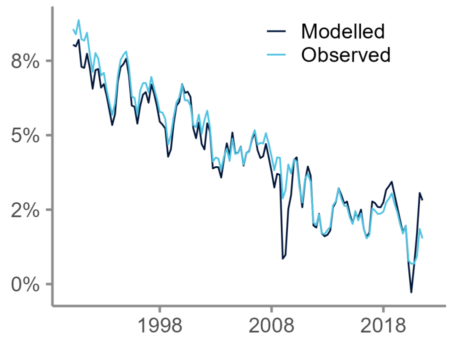

(D) 30-year

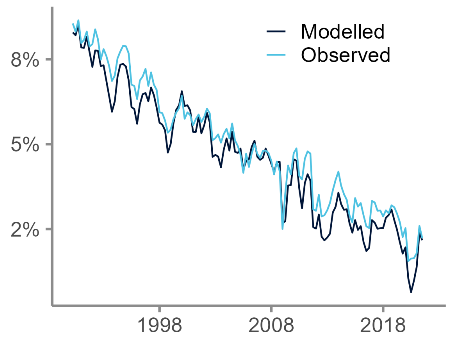

NOTE: The figures shows the historical fit to nominal US Treasury bonds.

Appendix B: Real fixed income return distributions

Figure 18 shows real returns on short- and long-duration fixed income portfolios. The 20-year distributions are narrower compared to the nominal return cases, in particular for the 1-year bond portfolio. This reflects the lower exposure of short-duration bonds to unexpected inflation. The comparisons of short- and long-duration portfolios are otherwise aligned with the results in the main text.

FIGURE 18 Distribution of cumulative real returns on 1- and 10-year government bonds at 3- and 20-year horizons

(A) 3-year horizon - Baseline

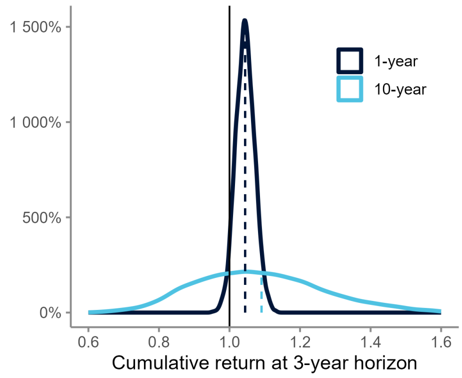

(B) 20-year horizon - Baseline

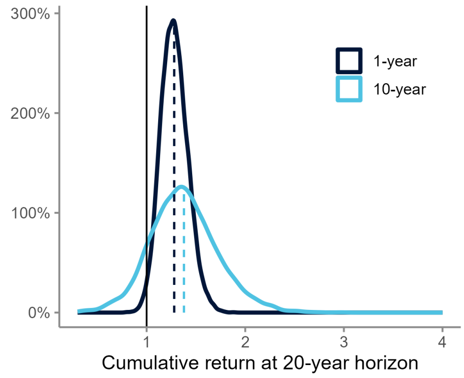

(C) 3-year horizon - Normalisation

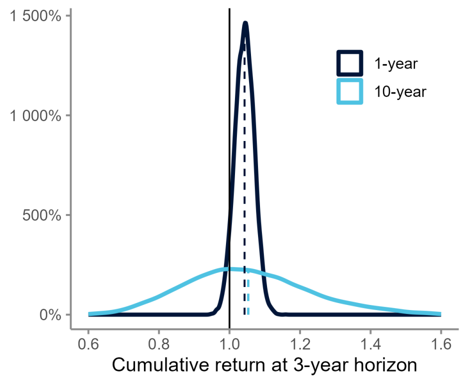

(D) 20-year horizon - Normalisation

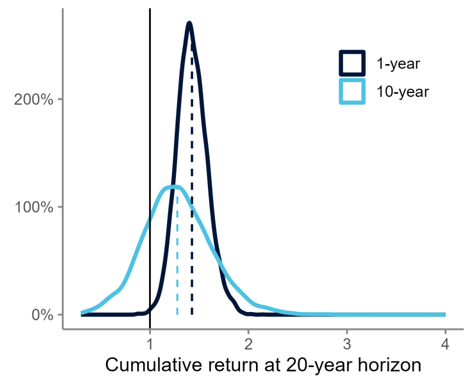

(E) 3-year horizon - Low Growth

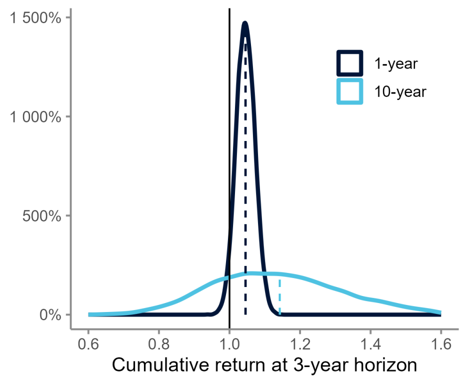

(F) 20-year horizon - Low Growth

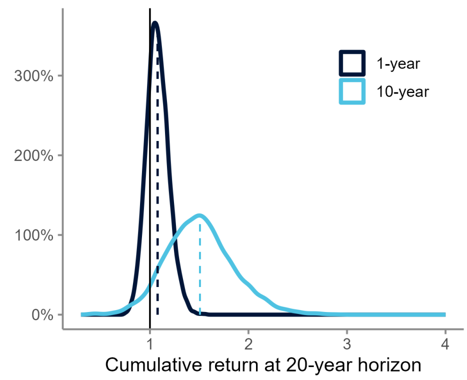

NOTE: Dashed lines indicate mean values of distributions.

Appendix C: Return distributions - percentiles

TABLE 3 Percentiles of cumulative nominal bond returns, Baseline calibration

Horizon (years)	1%	5%	25%	50%	75%	95%	99%
			Panel A. 1-year bond				
1	1.03	1.03	1.03	1.04	1.04	1.05	1.05
3	1.05	1.07	1.09	1.11	1.13	1.16	1.18
5	1.07	1.10	1.15	1.18	1.22	1.27	1.31
10	1.12	1.20	1.30	1.39	1.48	1.63	1.76
20	1.31	1.45	1.72	1.94	2.24	2.87	3.74
			Panel B. 10-year bond				
1	0.81	0.88	0.98	1.06	1.14	1.28	1.39
3	0.80	0.89	1.03	1.14	1.27	1.48	1.65
5	0.83	0.93	1.09	1.21	1.36	1.59	1.77
10	0.92	1.07	1.27	1.43	1.61	1.89	2.11
20	1.23	1.44	1.79	2.08	2.40	2.99	3.56

TABLE 4 Percentiles of cumulative nominal bond returns, Low Growth calibration

Horizon (years)	1%	5%	25%	50%	75%	95%	99%
			Panel A. 1-year bond				
1	1.03	1.03	1.03	1.04	1.04	1.05	1.05
3	1.04	1.06	1.09	1.11	1.13	1.15	1.17
5	1.05	1.08	1.13	1.16	1.20	1.25	1.29
10	1.05	1.11	1.21	1.28	1.37	1.50	1.63
20	0.99	1.11	1.30	1.45	1.65	2.05	2.55
			Panel B. 10-year bond				
1	0.81	0.88	0.99	1.07	1.15	1.29	1.40
3	0.84	0.93	1.08	1.19	1.32	1.53	1.70
5	0.89	0.99	1.17	1.30	1.45	1.71	1.90
10	1.04	1.19	1.41	1.58	1.77	2.08	2.32
20	1.23	1.43	1.76	2.01	2.33	2.87	3.36

TABLE 5 Percentiles of cumulative nominal bond returns, Normalisation calibration

Horizon (years)	1%	5%	25%	50%	75%	95%	99%
			Panel A. 1-year bond				
1	1.03	1.03	1.03	1.04	1.04	1.05	1.05
3	1.05	1.07	1.09	1.11	1.13	1.16	1.18
5	1.08	1.11	1.16	1.19	1.23	1.28	1.32
10	1.19	1.26	1.37	1.46	1.56	1.73	1.89
20	1.57	1.75	2.08	2.37	2.76	3.63	4.98
			Panel B. 10-year bond				
1	0.80	0.87	0.97	1.05	1.14	1.27	1.37
3	0.77	0.86	1.00	1.11	1.23	1.43	1.58
5	0.79	0.88	1.04	1.16	1.30	1.52	1.70
10	0.84	0.99	1.19	1.34	1.51	1.78	2.01
20	1.20	1.44	1.80	2.10	2.44	3.08	3.58

TABLE 6 Percentiles of cumulative returns on a 70-30% equity/bond portfolio, Negative bond-equity correlation

Horizon (years)	1%	5%	25%	50%	75%	95%	99%
			Panel A. Equity and 1-year bond				
1	0.90	0.94	1.01	1.07	1.13	1.23	1.31
3	0.93	1.00	1.13	1.23	1.35	1.56	1.72
5	0.98	1.09	1.26	1.40	1.56	1.81	2.04
10	1.26	1.41	1.68	1.90	2.17	2.64	3.07
20	2.15	2.48	3.09	3.68	4.42	6.03	8.00
			Panel B. Equity and 10-year bond				
1	0.93	0.97	1.03	1.08	1.13	1.21	1.28
3	0.99	1.06	1.17	1.26	1.36	1.53	1.67
5	1.08	1.17	1.33	1.44	1.58	1.80	1.97
10	1.43	1.57	1.82	2.01	2.24	2.62	2.93
20	2.58	2.93	3.55	4.13	4.83	6.22	7.80

TABLE 7 Percentiles of cumulative returns on a 70-30% equity/bond portfolio, Positive bond-equity correlation

Horizon (years)	1%	5%	25%	50%	75%	95%	99%
			Panel A. Equity and 1-year bond				
1	0.88	0.92	1.01	1.07	1.14	1.24	1.32
3	0.91	1.00	1.12	1.22	1.33	1.53	1.67
5	1.00	1.09	1.25	1.38	1.52	1.75	1.94
10	1.27	1.41	1.66	1.86	2.09	2.49	2.84
20	2.19	2.47	3.03	3.55	4.21	5.66	7.55
			Panel B. Equity and 10-year bond				
1	0.85	0.92	1.01	1.08	1.15	1.27	1.36
3	0.90	0.99	1.13	1.24	1.36	1.57	1.72
5	1.00	1.10	1.27	1.40	1.55	1.80	1.99
10	1.31	1.46	1.71	1.91	2.14	2.52	2.83
20	2.33	2.66	3.24	3.77	4.41	5.75	7.15

Appendix D: Zero average term premium

The term premium is unobserved and needs to be estimated, and there is uncertainty around its average level. While there are theoretical arguments for why the term premium should be positive on average, we should consider the possibility that it is not far from zero (e.g. government bonds are special/flight to safety). We therefore explore a scenario where the average term premium is zero. We keep other parameters unchanged relative to the Baseline calibration. Figures 19 and 20 compare short- and long-duration under this scenario. While the long-duration returns are marginally lower on average than under the Baseline calibration, the effect appears to be relatively muted.

FIGURE 19 Distribution of cumulative returns on 1- and 10-year nominal government bonds at the 20-year horizon – zero average term premium.

(A) 1-year bond **(B)** 10-year bond

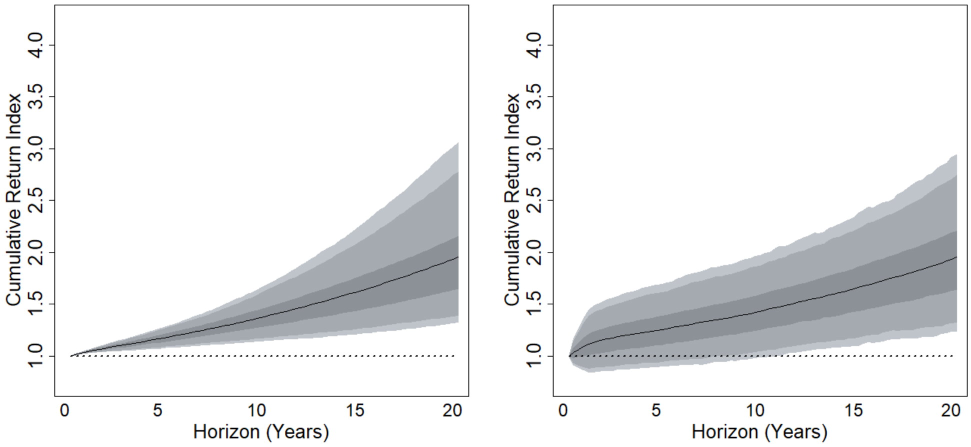

NOTE: Dark band covers 50 percent probability, medium band 90 percent, and light band 95 percent. Black line shows the expected value each quarter.

FIGURE 20 Distribution of cumulative returns on 1- and 10-year nominal government bonds at 3- and 20-year horizons – zero average term premium.

(A) 3-year horizon **(B)** 20-year horizon

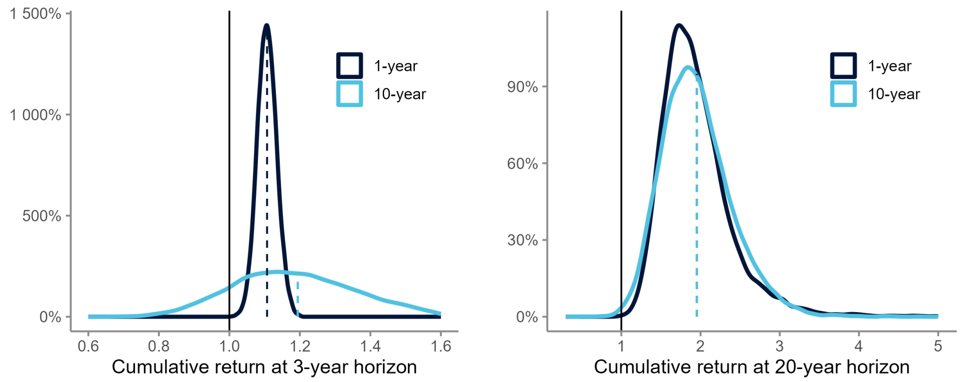

NOTE: Dashed lines indicate mean values of distributions.