Interest Rates and Equity Valuations*

Niels Joachim Gormsen and Eben Lazarus

June 2025

Abstract

The literature often seeks to determine the effect of interest rates on equity valuations, but both are endogenous and their comovement depends on the structural drivers underlying interest-rate changes. We show that changes in real rates can come from changes in expected growth, risk, or "pure discounting." We characterize the effect on equity valuations for each of the three shocks, and show that only pure discount rate shocks are transmitted one for one to equity valuations, with little or negative transmission of growth and risk shocks. Implementing our decomposition with a global panel of growth expectations and asset prices, we find: (i) a weak unconditional relation between stock valuations and real rates, but (ii) a strong relation between valuations and the pure discounting component of rates, with pure discount rate changes explaining over 80% of the cross-country changes in stock valuations since 1990. In the U.S. data, we find that 35% of the decline in interest rates is attributable to the pure discounting term, implying that only a fraction of the change in rates has passed through directly to equities. We also use our decomposition to speak to higher-frequency returns; explain interest-rate exposures in the cross-section of stocks; estimate a sizable duration-matched equity premium; and unpack the effects of policy-induced interest-rate shocks.

Keywords: Stock prices, interest rates, duration, long-term growth

JEL Codes: G10, G12, E44, F30

*An early draft of this paper circulated under the title "Equity Duration and Interest Rates." We are grateful to Michael Bauer, Martin Beraja, Jules van Binsbergen, Darrell Duffie, Matthieu Gomez, Rohan Kekre, Amir Kermani, Ralph Koijen, Arvind Krishnamurthy, Martin Lettau, Dmitry Livdan, Hanno Lustig, Peter Maxted, Thomas Mertens, Emi Nakamura, Lasse Pedersen, David Sraer, Jón Steinsson, and seminar participants at the SF Fed, UC Berkeley, and Stanford GSB for very helpful comments. Gormsen thanks the Fama-Miller Center, the Fama Faculty Fellowship, and the Center for Big Data in Finance (grant no. DNRF167) for financial support. Gormsen is at Copenhagen Business School, University of Chicago Booth School of Business, NBER, and the Danish Finance Institute, and Lazarus is at the Haas School of Business, University of California, Berkeley. Contact: niels.gormsen@chicagobooth.edu and lazarus@berkeley.edu.

1. Introduction

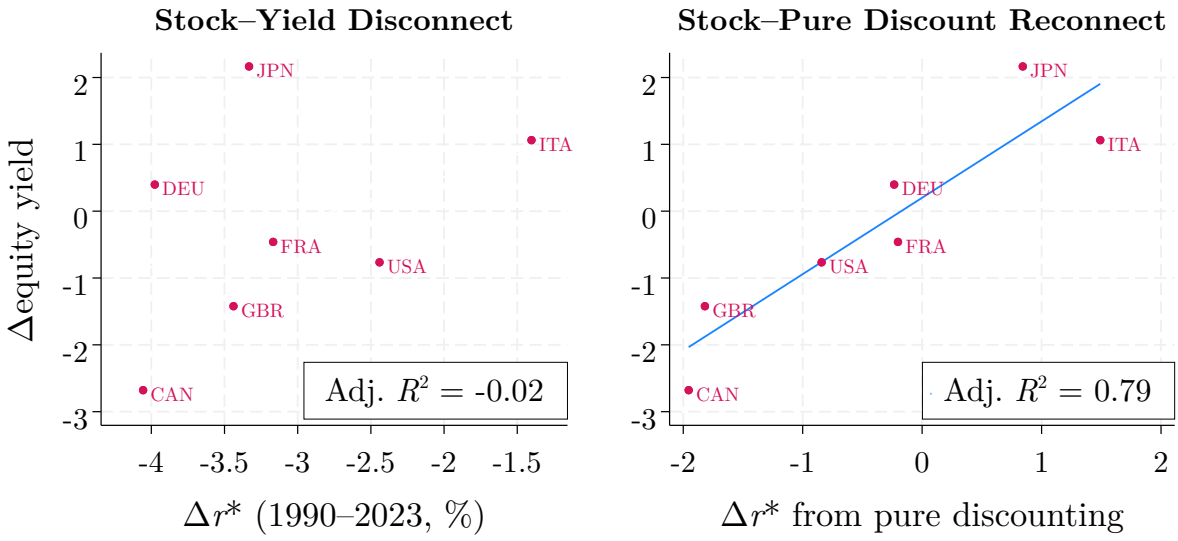
Advanced economies' long-term interest rates have declined significantly in recent decades. How do such changes in rates transmit to equity valuations? A potentially tempting line of reasoning is to assume that equity discount rates move one for one with interest rates; given that equity is a long-duration claim, this assumption then implies that stock valuations should have increased significantly as a result of the secular decline in interest rates. Casual evidence might appear consistent with this view: in the U.S., for example, the market's equity yield has declined substantially in recent decades, corresponding to a large increase in equity valuations over this period.

Empirically, however, there is no clear relationship between long-term changes in equity valuations and interest rates. The left panel of Figure 1 presents one view of this stock–yield disconnect: across G7 economies, the change in a country's equity yield since 1990 is effectively completely unrelated to that country's change in the trend long-term real rate r^*, described further below. In addition to the example in Figure 1, it is well known that the correlation between stock and bond returns is weak and often negative — another example of the stock-yield disconnect.

The apparent stock-yield disconnect arises because the interest-rate sensitivity of stock prices is more complicated than alluded to in our opening paragraph. Interest rates are determined endogenously and may decline for multiple possible structural reasons, each of which may affect equity differently. Interest rates may, for instance, decline because of a decrease in expected growth rates in the economy, which — keeping all else constant — will decrease equity prices and thus mute the effects of the decline in interest rates on equity prices.

In this paper, we provide a framework and measurement approach to control for the underlying drivers of interest-rate movements and estimate the interest-rate sensitivity of equity prices. We start with a simple but general theoretical decomposition under which any change in trend real rates can be split into three mutually exclusive underlying shocks: a change in expected growth, a change in uncertainty, and a pure discounting shock akin to a change in the rate of time preference. We then characterize how these shocks transmit to equity valuations. Only the pure discounting shock transmits one for one from rates to equity yields, thereby inducing perfect comovement between stocks and duration-matched bonds. The remaining terms, by contrast, induce a weak and theoretically ambiguous relationship between stocks and bonds: a growth-rate shock affects both equity discount rates and cash-flow growth, while an uncertainty shock causes interest rates and equity risk premia to move in opposing directions. Isolating the pure discounting component of real rates is therefore key

Figure 1: Preview of Main Results: Long-Term Decomposition



Notes: This figure plots the country-level changes in equity yields against changes in trend interest rates (left panel) and in the estimated change in the pure discounting component of interest rates (right panel), for G7 economies. The sample is 1990–2023, or the longest available span for the given country. Details on the measurement of the equity yield, trend long-term interest rate r^*, and the pure discounting component of the change in r^* are provided in Section 3.2.

for understanding how much any given change in interest rates passed through to equities.

We next implement the decomposition empirically using a combination of survey data and option prices to estimate the underlying structural drivers. The punchline of our empirical implementation is that most of the secular changes in stock prices over the past 35 years can be explained by movements in the pure discounting part of interest rates. This result is illustrated in the right panel of Figure 1, which shows that 80% of the changes in valuation ratios of equities in G7 countries can be explained by changes in the pure discounting part of interest rates. This pure discounting component is estimated purely from our decomposition for interest rates, so there is nothing mechanical about the tight fit in explaining almost the entirety of the country-by-country change in equity valuations over recent decades.

The pure discounting part of interest rates also explains a substantial share of fluctuations in stock prices at higher frequencies — consistent with the long cash-flow duration of the stock market — and the pure discount term can be used to understand the pricing of the cross-section of equities. Our framework allows us to revisit outstanding puzzles in the literature and estimate the channel through which monetary policy influences stock prices, all of which we elaborate on below.

The key input for our measurement is an international panel of long-term professional forecasts for interest rates, inflation, and growth rates, which we obtain from Consensus

Economics. We back out forecast-implied series for trend real rates r^* and trend growth rates by country, and we augment these with option-based measures of uncertainty to estimate our interest-rate decomposition. After stripping out growth-rate and uncertainty changes, the remaining interest-rate change is our estimate of the pure discounting shock.

In the U.S., we attribute around 35% of the decline in r^* since 1990 to pure discount-rate changes, and the remaining 65% to the other components. So while equities have benefitted somewhat from the decline in U.S. interest rates, assuming full pass-through of r^* to equity yields would overstate the effect by close to three times. And the passthrough of the decline in rates to equities has been even lower in most other G7 economies. This partial pass-through of interest rates to equity valuations speaks to a wide range of questions studied in recent literature; as discussed in the literature review just below, we apply our results to better understand, for example, the long-term performance of stocks versus bonds, and the degree to which changes in household portfolio values reflect purely "paper" gains.

After considering the long-term trends, we then apply our decomposition to consider the drivers of interest-rate changes at higher frequencies and their effects on equity returns. Without adjusting for the endogeneity of interest rates, the raw relation between market returns and yield changes is small and imprecisely estimated. But when we implement our decomposition, pure discount-rate shocks generate strong negative comovement between Δr_t^* and annual equity returns. The loading of stock returns on pure discounting shocks provides a theoretically well-founded measure of equity duration, and we estimate a market duration of about 20 years in the U.S. data.[1] By contrast, equity returns have a small and insignificant relation to the interest-rate change attributable to changing expected growth rates, and a positive relation with the interest-rate change attributable to uncertainty shocks. These offsetting components illustrate why equity duration is not equivalent to the price sensitivity to arbitrary changes in interest rates, and why one must isolate the pure discounting component to estimate duration.

Similarly, we find using forecasting regressions that long-term yields by themselves do not predict future equity returns, providing further evidence that risk premia often comove negatively with yields. The pure discounting term, by contrast, strongly and significantly predicts returns, providing further evidence that it strips out confounding shocks to yields. Finally, we conduct a return accounting exercise to estimate the contribution of each of the three interest-rate shocks to market-level equity returns on a rolling basis. Pure discounting decreases are important for explaining the strong performance of U.S. stocks and bonds in

[1]We consider this preliminary estimate to be a likely lower bound for the true duration, given the measurement uncertainty in our pure discounting shock. This is nonetheless a large and significant estimate for equity duration, and it has the benefit of being an ex ante measure that does not require estimating equity cash-flow growth rates or discount rates in realized historical data.

the 1990s. But the realized shocks to the pure discounting term (including a positive shock post–2020) have roughly offset over the period since 2000, generating a roughly zero net effect of such discounting changes on interest rates or equities over the most recent decades.

We next use our decomposition to better understand the cross-section of stocks and their exposure to interest rates. Following Gormsen and Lazarus (2023), we sort firms by their predicted cash-flow duration and measure these duration-sorted portfolios' returns. We show that these portfolios do not differ in their exposure to raw interest-rate changes, but that the long-duration firms have significantly greater exposure to the pure discounting shock. This holds in spite of the unconditional negative alpha to long-duration relative to short-duration stocks, and it implies a sizable spread (greater than 20 years) in the duration of long- versus short-duration firms' cash flows. These results provide a further out-of-sample validation of both the duration sort and of the construction of the pure discounting term.

All of our baseline results operate under the assumption that expected output growth (as is relevant for risk-free rates) is proportional to expected dividend growth (as is additionally relevant for equity valuations). We therefore provide a set of theoretical and empirical robustness results on the effects of changes in the profit share of income, or the ratio of earnings to aggregate output. Greenwald, Lettau, and Ludvigson (2025) estimate that about 40% of U.S. equity returns since 1989 are attributable to unanticipated increases in the profit share. As they note, such shocks increase both prices and cash flows; in our context, the effect of profit-share shocks on equity yields depends on the resulting changes in the path of expected *future* cash-flow growth. Using separate forecast data on earnings and dividends, we find that our main empirical results in the U.S. data are close to unchanged even when allowing for separate output growth and cash-flow growth processes. Changing profit shares thus seem to affect equity prices mainly through contemporaneous (i.e., already materialized) cash flows, rather than expected future growth rates.

Taken together, our results show that isolating the pure discounting term is essential for understanding how shocks to interest rates are reflected in stock prices. While changes to this term are equivalent to changes in the pure rate of time preference for the marginal investor, we do not view this as the only (or main) source of likely variation. To better interpret this term, we provide further results showing how our decomposition should be understood when households can freely invest in other countries' bonds or stocks. Net capital inflows can induce a decline in the pure discounting component of interest rates — that is, a decrease in interest rates that is not accounted for by a decline in expected growth or increased uncertainty — and such flows accordingly represent a candidate source of variation for the pure discounting term. We then show how the changes in the estimated pure discounting terms align well with cross-country capital flows empirically.

Implications and Connections to Recent Literature

Our characterization of the pass-through of interest rates to equities speaks to a range of findings and questions raised in recent literature:

(i) van Binsbergen (2024) shows that long-term bond portfolios have performed nearly as well as equities in recent decades.[2] Our results provide an explanation for this finding: much of the decline in interest rates has arisen from growth-rate and uncertainty shocks that should not increase equity valuations. In our setting, the natural duration-matched benchmark for equity returns is a pure discounting claim. We conduct an additional analysis measuring the returns on such a claim in the data, and we estimate a large and stable duration-matched equity premium.

(ii) Numerous papers have studied the degree to which declining interest rates have affected the value of different households' overall portfolios, including equities and other risky assets (see, for example, Catherine, Miller, Paron, and Sarin, 2023, and Greenwald, Leombroni, Lustig, and Van Nieuwerburgh, 2023). If r^* declines have passed through fully to these risky assets (which are held disproportionately by wealthy households), then much of the increase in wealth inequality in recent decades may reflect purely "paper" gains.[3] Our findings help speak to this debate. We estimate that a sizable share of the decline in r^* did not transmit directly to equity valuations, indicating that much of the resulting increase in portfolio valuations (and thereby inequality) was non-mechanical.

(iii) Recent work has also considered how factor returns and cross-sectional anomalies may have been affected by the change in interest rates. van Binsbergen, Ma, and Schwert (2024) argue that anomaly portfolios investing in firms with short-duration cash flows (such as value, or high-book-to-market, portfolios) would have exhibited better performance in recent decades in a counterfactual without the large decline in interest rates. Maloney and Moskowitz (2021) emphasize a weak relation between interest rates and value stocks' underperformance in recent years, which challenges theories of the value premium relating to cash-flow duration.[4] Our framework and

[2]Similarly, Andrews and Gonçalves (2020) estimate a near-zero risk premium on long-maturity dividends relative to long-maturity bonds.

[3]While Fagereng, Gomez, Gouin-Bonenfant, Holm, Moll, and Natvik (2024) do not advance this argument, they provide a succinct review of literature making this claim.

[4]Maloney and Moskowitz (2021) state on page 4, "Dechow, Sloan, and Soliman (2004), Lettau and Wachter (2007), and Gormsen and Lazarus ([2023]) characterize value stocks as low-duration assets with near-term cash flows and growth stocks as high-duration assets, such that a long-short value strategy is a negative-duration asset that is sensitive to falling interest rates. This story implies that falling bond yields from 2010 to 2020 acted as a strong tailwind for growth stocks and a headwind for value stocks, driving

results show that the exposure of a given portfolio to unadjusted interest-rate changes depend on the underlying driver of the change in interest rates. When focusing on the pure discounting component, we find that value stocks indeed have less exposure to this pure discounting shock than growth stocks. We then use this to clarify the role of interest-rate declines on the poor performance of value stocks, finding that the pure discounting component explains some, but not nearly all, of the underperformance of value in recent years.

(iv) As a further application relevant for both asset pricing and macroeconomics, we use our decomposition and estimation results to help unpack the effects of surprise changes in short-term interest rates by monetary policymakers. While some papers have treated the resulting changes in long-term rates as if they represent pure discounting shocks, this is not necessarily a valid assumption: while the change in the short-term rate is indeed exogenous, the long-term yield change depends on changes to the pure discount rate *as well as* changes to the market's perceived long-term growth and uncertainty. We use our main estimation results for stock returns and yield changes, along with high-frequency asset-price changes observed around monetary policy announcements, to back out announcement-specific changes in both the pure discounting term and expected growth rates. On average, we find that most of the change in long-term yields around policy shocks indeed stems from the pure discounting component. But we find evidence as well that growth-rate expectations change in a manner consistent with an information effect (Nakamura and Steinsson, 2018) on average, with meaningful announcement-specific heterogeneity in this response.

Overall, our results provide a new toolkit with which to understand how changes in interest rates have — and have not — affected a range of risky assets and aggregate outcomes both over the short and long run.

Additional Related Literature

In addition to the tie-ins to recent papers described above, our paper relates to a long literature analyzing the time-varying relationship between stocks and bonds. Much of this work has focused on the comovement between stocks and *nominal* bonds, and much of it has focused on characterizing higher-frequency (e.g., daily to quarterly) comovements. David and Veronesi (2013) and Campbell, Pflueger, and Viceira (2020) estimate model-implied sources of the shift in the stock–bond correlation from positive to negative in the early 2000s,

value-tilted portfolio returns lower."

attributing much of the shift to changes in inflation risk dynamics.[5] Our focus on real interest rates, and our use of a general decomposition as opposed to a fully parameterized model, distinguish us from these and related papers. Baele, Bekaert, and Inghelbrecht (2010) use a dynamic factor model to characterize the drivers of this changing correlation, arguing that non-fundamental factors are important. Chernov, Lochstoer, and Song (2023) argue for a real channel in which the relative importance of permanent vs. transitory consumption shocks drives the comovement, while Laarits (2022) argues for changing uncertainty and a precautionary savings channel. Our focus on longer-term secular drivers of stock–bond comovements distinguish us from this work.

Our paper is somewhat closer to work studying such long-term comovement more directly. Campbell and Ammer (1993) use a vector autoregression to characterize both short-term stock–bond comovements and how these relate to longer-term expected returns. We differ in both methodological approach (we use surveys and asset prices, rather than a VAR, to estimate forward-looking expectations) and in the framework taken to the data: we aim to measure the relation of bond and equity valuations and returns to three fundamental variables underlying interest rates, rather than relating these values to future expected returns and cash flows for each of the two assets. Our framework is somewhat more related conceptually to a simpler two-period framework considered by Barsky (1989), but with fewer parametric restrictions and with pure discounting shocks playing a key role in our case. Barsky's contribution is, in addition, theoretical rather than empirical. Farhi and Gourio (2018) estimate a simple neoclassical growth model with markups and intangibles to account for secular changes in interest rates and risky asset valuations. The long-run questions they pose, and the simple framework taken to the data, overlap in spirit with ours. We also find evidence consistent with their view that risk-free rates have comoved negatively with equity risk premia in recent decades. Our decomposition and estimation exercise, however, is different — and somewhat less tightly parameterized — than theirs.[6]

In addition to these two main literatures related to the comovement of stocks and bonds, our framework relates to recent work using different assets' comovements to distinguish which channels are most important for long-term price variation. Much of this work uses comovements between exchange rates, relative interest rates, and fundamentals to understand exchange-rate determination; recent examples include Lustig and Verdelhan (2019), Itskhoki

[5] See also Song (2017), and see David and Veronesi (2016) for a review.

[6] Further related literature along these lines includes work on the long-term drivers of stock prices. As a prominent recent example, Greenwald, Lettau, and Ludvigson (2025) argue that increased profit shares are a key driver of equity-price increases in the U.S. data. While their model-based estimation differs from ours, our results appear consistent with theirs, as discussed above and further in Section 4.1. And like us, they estimate that less than half of the increase in U.S. equity prices is attributable to interest-rate changes.

and Mukhin (2021), Jiang, Krishnamurthy, and Lustig (2024), and Kekre and Lenel (2024).[7] Kekre and Lenel's setting and framework provide a particularly useful contrast with ours. They measure the importance of demand shocks for long-term exchange-rate determination. But as they note, such demand shocks encompass both time-preference shocks and growth-rate shocks, and exchange rates and interest rates by themselves do not allow one to distinguish these two sources of variation. Considering equity prices in addition to interest rates, as we do, allows one to discriminate between these two different shocks. We find an important role for both in explaining long-term variation in real rates and equity valuations.

Organization

We begin with our theoretical decompositions in Section 2. We then turn to our data, measurement approach, and main findings in Section 3, and Section 4 provides a set of robustness results. In Section 5, we analyze additional implications of our findings for asset prices in recent decades and related literature. Section 6 discusses and concludes. Derivations and additional results can be found in the Appendix.

2. Theoretical Decompositions

This section provides our theoretical decomposition for the trend real rate. There are of course arbitrarily many valid decompositions for real rates. Our goal is a decomposition of this endogenous object into interpretable fundamental components, in a manner that both (i) is empirically estimable and (ii) contains one component that induces perfect comovement of bonds and stocks. Isolating this last component will then allow us to measure the degree to which interest-rate changes transmit to equity valuations.

We begin in Section 2.1 with a general decomposition with minimal assumptions on the fundamentals or the stochastic discount factor. We then specialize to a more interpretable consumption-based version of the decomposition in Section 2.2. We consider what each term in the real-rate decomposition means for equity prices in Section 2.3, and for equity duration in Section 2.4.

[7] Others, including Pavlova and Rigobon (2007), Camanho, Hau, and Rey (2022), Atkeson, Heathcote, and Perri (2024), and Auclert et al. (2024), additionally consider asset valuations and capital flows. We briefly analyze exchange rates and capital flows in Section 4.2. Kremens, Martin, and Varela (2024), meanwhile, show that fundamentals are closely tied to survey forecasts of long-horizon currency appreciation, which are themselves strong predictors of actual currency changes; this finding of a link between fundamentals, surveys, and future returns mirrors ours in a different market.

2.1 A General SDF-Based Version

We start with a general stochastic discount factor (SDF) M_{t+1} such that $\mathbb{E}_t[M_{t+1}R_{t+1}] = 1$ for an arbitrary asset's gross return R_{t+1}. This implies $R^f_{t+1} = 1/\mathbb{E}_t[M_{t+1}]$, where R^f_{t+1} is the real risk-free rate. Taking logs (and denoting logged variables in lowercase),

$$r^f_{t+1} = -\mathbb{E}_t[m_{t+1}] - L_t(M_{t+1}), \tag{1}$$

where $L_t(M_{t+1}) \equiv \log \mathbb{E}_t[M_{t+1}] - \mathbb{E}_t[m_{t+1}]$ is the conditional entropy of the SDF.[8]

For now, we put very little structure on the SDF. We assume that the log SDF can be additively decomposed as follows:

$$m_{t+1} = \underbrace{-\rho_t}_{\substack{\text{predetermined} \\ \text{trend}}} - \underbrace{(f(X_{t+1}) - f(X_t))}_{\substack{\text{difference for} \\ \text{Markov } X}} + \underbrace{\varepsilon_{t+1}}_{\substack{\text{mean 0} \\ \text{martingale diff.}}}. \tag{2}$$

This representation is based on an additive decomposition constructed following Hansen (2012, Theorems 3.1–3.2), and it holds under a general set of primitive assumptions. See Appendix A.1 for formal details and a discussion. As discussed in the appendix, the term $f(X_{t+1}) - f(X_t)$ is either stationary or difference-stationary.

In interpreting (2), the trend $-\rho_t$ shifts the intertemporal marginal rate of substitution m_{t+1} in all states, so ρ_t can be thought of as a time discount rate. We interpret the Markov state X_{t+1} as determining aggregate cash flows, so $f(X_{t+1}) - f(X_t)$ can be thought of as the realized marginal utility from cash flow growth. Finally, ε_{t+1} is the remaining martingale component of the log SDF. These terms' interpretation will map to their interpretation in the consumption-based framework in the next subsection.

Plugging (2) into (1), the log risk-free rate satisfies

$$r^f_{t+1} = \underbrace{\rho_t}_{\text{trend (discounting)}} + \underbrace{\mathbb{E}_t[f(X_{t+1}) - f(X_t)]}_{\text{expected growth}} - \underbrace{L_t(M_{t+1})}_{\text{uncertainty/prec. savings}}. \tag{3}$$

The first two terms' labels align with the interpretations discussed above. For the labeling of $L_t(M_{t+1})$ as an uncertainty or precautionary savings term, note that by definition of entropy,

$$L_t(M_{t+1}) = \sum_{n=2}^{\infty} \frac{\kappa_{n,t}(m_{t+1})}{n!}, \tag{4}$$

[8] This is also the starting point for studying exchange-rate puzzles in Backus, Foresi, and Telmer (2001), Hassan, Mertens, and Wang (2024), and Jiang, Krishnamurthy, and Lustig (2024), and for studying disasters and risk premia in Backus, Chernov, and Martin (2011).

where $\kappa_{n,t}(m_{t+1})$ is the n^{th} conditional cumulant of the log SDF distribution (assumed to be finite for all n). Conditional entropy therefore encodes the higher ($n \geqslant 2$) moments of marginal utility, as is standard.

Our main interest will be in understanding changes in the trend risk-free rate r_t^*. Analogous to Bauer and Rudebusch (2020), we define this as the Beveridge–Nelson trend in the one-period real rate, $r_t^* \equiv \lim_{\tau\to\infty} \mathbb{E}_t[r^f_{t+\tau+1}]$. But unlike Bauer and Rudebusch, we are interested in somewhat longer-horizon real rates: we think of one period as being equal to the cash-flow duration of the overall equity market. As a result, we will not directly consider the term premium embedded in these longer-term rates. Considering a long period length is equivalent to considering a multi-period zero-coupon yield directly in (1): $R^f_{t,t+\tau} = 1/\mathbb{E}_t[M_{t,t+\tau}]$ for any $\tau \geqslant 1$, where $M_{t,t+\tau} = M_{t+1} \cdots M_{t+\tau}$, so (1) applies when replacing "$t+1$" with "$t, t+\tau$." Term premia affect interest rates through the entropy term in (3), as studied by Backus, Boyarchenko, and Chernov (2018).

For the trend real rate r_t^*, equation (3) directly implies that it satisfies

$$r_t^* = \rho_t^* + \widetilde{g}_t^* - L_{t,M}^*, \tag{5}$$

where $\rho_t^* = \rho_t$, $\widetilde{g}_t^* = \lim_{\tau\to\infty} \mathbb{E}_t[f(X_{t+\tau+1}) - f(X_{t+\tau})]$ and $L_{t,M}^* = \lim_{\tau\to\infty} \mathbb{E}_t[L_{t+\tau}(M_{t+\tau+1})]$. This is the first version of our real-rate decomposition into three terms corresponding to discounting, expected growth, and uncertainty. The analysis in this section showed that such a decomposition can be derived quite generally — up to the issue of interpretation of each of the three terms — starting from an additive decomposition of the log SDF.

2.2 A Consumption-Based Version

To put more structure on the decomposition in (5), we now consider a more standard consumption-based framework. We assume an endowment economy in which a representative agent has power utility over consumption,

$$U_t = \mathbb{E}_t \sum_{\tau=0}^{\infty} \beta_t^\tau \frac{C_{t+\tau}^{1-\gamma}}{1-\gamma}. \tag{6}$$

The time discount factor β_t and corresponding rate of time preference $\rho_t = -\log \beta_t$ are potentially time-varying. We assume that relative risk aversion γ, or the inverse elasticity of intertemporal substitution, is constant.[9]

[9]See Appendices A.1–A.2 as well for extensions with Epstein–Zin preferences, time-varying risk aversion, or other departures from the basic model. As discussed there, decompositions of the form (5) or (8) still hold in these alternative specifications; see equation (A.18) for an exact solution in a tractable Epstein–Zin case.

Given (6), the log SDF is $m_{t+1} = -\rho_t - \gamma g_{t+1}$, where $g_{t+1} \equiv c_{t+1} - c_t$ is log consumption growth. Plugging this into (1),

$$\begin{aligned} r^f_{t+1} &= \rho_t + \gamma \mathbb{E}_t[g_{t+1}] - L_t(M_{t+1}) \\ &= \rho_t + \gamma \mathbb{E}_t[g_{t+1}] - \sum_{n=2}^{\infty} \frac{(-\gamma)^n \kappa_{n,t}(g_{t+1})}{n!}, \end{aligned} \tag{7}$$

where the final expression for $L_t(M_{t+1})$ in terms of the growth-rate cumulants $\kappa_{n,t}(g_{t+1})$ is as in Backus, Chernov, and Martin (2011) or Martin (2013); see Appendix A.2. In a lognormal setting, this simplifies to the familiar solution $r^f_{t+1} = \rho_t + \gamma \mathbb{E}_t[g_{t+1}] - \frac{\gamma^2}{2} \mathrm{Var}_t(g_{t+1})$.

Given (7), the trend real rate can be expressed as

$$r^*_t = \rho^*_t + \gamma g^*_t - L^*_{t,M}, \tag{8}$$

where $\rho^*_t = \lim_{\tau \to \infty} \mathbb{E}_t[\rho_{t+\tau}]$, $g^*_t = \lim_{\tau \to \infty} \mathbb{E}_t[g_{t+\tau+1}]$, and $L^*_{t,M} = \lim_{\tau \to \infty} \mathbb{E}_t[L_{t+\tau}(M_{t+\tau+1})]$.[10] Thus exactly as in (5), the real rate can move due to changes in (i) time preference (a stand-in for pure discounting shocks), (ii) expected growth rates (via an intertemporal substitution channel), or (iii) risk or uncertainty (via a precautionary savings channel).

2.3 Implications for Equity Prices

We now move to equity and ask how each of the three channels — pure discounting, expected growth, and uncertainty — transmit from risk-free rates to equity valuations. Doing so requires further structure on equity cash flows. Our framework here will be quite standard. We derive a version of a Gordon growth formula for equity dividend yields; while we do so in three slightly different environments, the basic pricing formulas will be quite similar (and, for the most part, familiar). The main point will be to show that the pass-through of each of the three components of our interest-rate decomposition to equity valuations applies in a simple, intuitive way, in a range of standard settings.

We start from the same consumption-based framework as in Section 2.2, though the main equity-valuation decompositions below in fact apply as well under Epstein–Zin utility (see Appendix A.3). For equity cash flows, we follow Campbell (1986) and Abel (1999) and model equity as a levered claim to consumption, paying dividends $D_t = C^\lambda_t$, with $\lambda > 0$. This imposes a tight link between dividend growth and output and consumption growth. We will see that this works well to explain the long-term trends observed in the expectations

[10] As discussed in Appendix A.1.3, in a stationary setting in which the limiting expectations here are constant, we redefine these terms as discounted sums of expected outcomes from date $t+1$ to ∞.

data, but this need not always be the case (particularly at high frequencies). We accordingly extend our analysis in Section 4.1 to allow for time variation in the profit share of output, which we then discipline in the data with survey expectations on profit growth. For now, however, we assume that λ is constant.

We denote the gross equity return by R^{mkt}_{t+1} and the log return by r^{mkt}_{t+1}, and define $\mu_t \equiv \mathbb{E}_t[r^{\text{mkt}}_{t+1}]$ and $rp_t \equiv \mu_t - r^f_{t+1}$. The equity yield is defined as

$$ey_t \equiv \log(1 + D_t/P_t), \tag{9}$$

where P_t is the price of the equity claim. This is slightly different from the usual log dividend–price ratio ($dp_t \equiv \log(D_t/P_t)$). We define ey_t as in (9) because it puts the equity yield in equivalent units as the log real rate. It also yields straightforward characterizations of steady-state ey^*_t building on results from Martin (2013) and Gao and Martin (2021).

We consider three cases for the dynamics of underlying fundamentals, each of which delivers approximately identical intuition for equity valuations.

Case I (Gordon Growth): Assume that log output growth $g_{t+1} = c_{t+1} - c_t$ is i.i.d. over time, with arbitrary distribution. We view this as a reasonable approximation given our focus on relatively long horizons (over which outcomes may be roughly conditionally i.i.d.), though this i.i.d. assumption will mean that any changes in moments or preference parameters are assumed to be unexpected permanent shocks. We write $g^*_t = \mathbb{E}_t[g_{t+1}] = \mathbb{E}_t[g_{t+\tau}]$ for all $\tau \geqslant 1$. Given i.i.d. g_{t+1}, the growth-rate cumulants $\kappa_{n,t+\tau}(g_{t+\tau+1})$ are also constant for all τ; we again include time indexes to allow for unanticipated shocks to these values. We similarly allow for unanticipated shocks to the preference parameter $\rho^*_t = \rho_t$.

This setting mirrors that of Martin (2013), and we apply and build on his results; see Appendix A.3 for details and derivations. Given the constant growth rates and discount rates, a Gordon growth formula applies as follows (where we use the i.i.d. assumption to set all relevant variables equal to their conditional steady-state values):

$$ey^*_t = r^*_t + rp^*_t - \lambda g^*_t. \tag{10}$$

The log equity premium satisfies

$$\begin{aligned} rp^*_t &= \sum_{n=2}^{\infty} \frac{\kappa_{n,t}(g_{t+1})}{n!} \left((-\gamma)^n - (\lambda - \gamma)^n\right) \\ &= L_t(M_{t+1}) - L_t(M_{t+1}R^{\text{mkt}}_{t+1}) = L^*_{t,M} - L^*_{t,MR}. \end{aligned} \tag{11}$$

As discussed in the appendix, the fact that $rp_t = L_t(M_{t+1}) - L_t(M_{t+1}R^{\text{mkt}}_{t+1})$ is fully general: it holds under no arbitrage and does not require i.i.d. fundamentals or any assumptions on utility. If $\lambda = \gamma$, then $M_{t+1}R^{\text{mkt}}_{t+1} = 1$, and $rp^*_t = L^*_{t,M}$.[11] Alternatively, for arbitrary λ and γ, if growth is lognormal so that $\kappa_{n,t}(g_{t+1}) = 0$ for $n > 2$, then $rp^*_t = \frac{1}{2}\lambda(2\gamma - \lambda)\text{Var}_t(g_{t+1}) = \frac{\lambda(2\gamma-\lambda)}{\gamma^2}L^*_{t,M}$.

Using (8), (10), and (11), we obtain a solution for equity yields summarized along with the real risk-free rate in the following result.

Result 1. *The steady-state real risk-free rate and equity dividend yield satisfy*

$$\begin{aligned} r^*_t &= \rho^*_t + \gamma g^*_t - L^*_{t,M}, \\ ey^*_t &= \rho^*_t + (\gamma - \lambda)g^*_t + (rp^*_t - L^*_{t,M}) \\ &= \rho^*_t + (\gamma - \lambda)g^*_t - L^*_{t,MR}. \end{aligned}$$

Changes in the risk-free rate can arise due to (i) pure discounting shocks (changes in ρ^*_t), (ii) growth-rate shocks (g^*_t), or (iii) risk (entropy) shocks ($L^*_{t,M}$). Each of the three has different implications for equity valuations:

(i) ***Pure discounting shocks:*** *Bonds and equity co-move perfectly, with ey^*_t increasing by 1 basis point for each 1 basis point increase in r^*_t.*

(ii) ***Growth-rate shocks:*** *Equity yields change by $\frac{\gamma-\lambda}{\gamma}$ per unit increase in r^*_t. If $\gamma = \lambda$ (e.g., with log utility and an unlevered consumption claim), ey^*_t is unaffected by changes in r^*_t induced by growth shocks. If $\lambda > \gamma$, growth shocks induce negative comovement.*

(iii) ***Risk shocks:*** *Equity yields change by $-\frac{\partial rp^*_t}{\partial L^*_{t,M}} + 1$ per unit increase in r^*_t if $\frac{\partial rp^*_t}{\partial L^*_{t,M}}$ is well-defined. Otherwise, equity yields change on average by $-\beta_L + 1$ per unit increase in r^*_t, where $\beta_L \equiv \frac{Cov(rp^*_t, L^*_{t,M})}{Var(L^*_{t,M})}$. If $\gamma = \lambda$, then ey^*_t is unaffected by changes in r^*_t induced by risk shocks. If $\beta_L > 1$, risk shocks induce negative comovement.*

The key implication of this result is that only the pure discounting channel generates perfect pass-through from interest rates to equity yields, and from bond prices to duration-matched stock prices (as discussed below). The range of past work assuming that the decline in interest rates has passed through to equity valuations — as discussed in Section 1 — has therefore implicitly assumed that the decline in r^*_t has arisen due to such pure discounting shocks. Changes in ρ^*_t can be thought of as capturing, for example, demographic changes, or something akin to a savings glut. We discuss such interpretations further in subsequent sections.

[11] This is a restatement of the fact that the Alvarez and Jermann (2005) SDF entropy lower bound holds with equality in the growth-optimal case (which corresponds to the current case in that $M_{t+1}R^{\text{mkt}}_{t+1} = 1$): $L_t(M_{t+1}) = \mathbb{E}_t[r^{\text{mkt}}_{t+1}] - r^f_{t+1}$.

For growth-rate changes, note that the equity yield depends on $r_t^* - \lambda g_t^*$ (our version of $r - g$), so a decline in g_t^* will have roughly offsetting effects since it decreases both discount rates and growth rates.[12] It is common to assume that $\gamma \leqslant \lambda$ — or, in the Epstein–Zin case, that $\frac{1}{\psi} \leqslant \lambda$, where ψ is the elasticity of intertemporal substitution — so that a decline in growth rates also decreases equity valuations (corresponding to a higher ey_t^*). This implies that growth-rate changes induce weakly negative comovement between bonds and stocks.

For changes in risk, there are offsetting effects on the risk-free rate and the risk premium. These changes may approximately offset or may cause stocks to move in the opposite direction of bonds. As $L_{t,M}^*$ loads on all the higher cumulants of the growth distribution as in (7), the stock-price response will depend on the specific parameter change underlying the risk shock. In Appendix A.3, we characterize the bond–stock comovement in three benchmark cases: (i) in a lognormal setting with power utility and $\gamma \neq \lambda$, r_t^* and ey_t^* comove positively given changes in risk, though the pass-through is less than one-for-one if $2\gamma > \lambda$; (ii) in a rare-disasters model as in Barro (2006), if $\gamma < \lambda$, then r_t^* and ey_t^* comove negatively given changes in the average disaster size (or, more generally, given changes in skewness or other odd moments); (iii) with Epstein–Zin utility, if $\gamma > 1$, $\psi > 1$, and $\lambda = 1$, then r_t^* and ey_t^* comove negatively given changes in even moments of the growth distribution (variance, kurtosis, and so on), as in Martin (2013). So while positive risk shocks robustly decrease the risk-free rate and generally induce a muted or negative stock–bond comovement, we treat their precise pass-through to stocks as an open question to be disciplined empirically.

Result 1 is our main decomposition for trend real rates and equity valuations. We now discuss how this decomposition carries through under more realistic assumptions on the evolution of fundamentals.

Case II (Drifting Steady State): Assume now that the distributions of growth rates and preference parameters are such that ey_t follows a martingale, $ey_t = \mathbb{E}_t[ey_{t+1}]$, similar to Campbell (2008). Campbell (2018) refers to this as a "drifting steady state" model for $ey_t^* = ey_t$, and Campbell and Thompson (2008) show that such a model has success at forecasting medium-to-long-horizon returns. To a first order for ey_{t+1} around its expectation ey_t, we have in this case that $ey_t = \mathbb{E}_t[r_{t+1}^{\text{mkt}} - \lambda g_{t+1}]$. (This follows Gao and Martin, 2021, and again see the appendix.) This implies that

$$ey_t^* = ey_t = \mathbb{E}_t[r_{t+1}^{\text{mkt}} - \lambda g_{t+1}] = \mathbb{E}_t[r_{t+2}^{\text{mkt}} - \lambda g_{t+2}] = \ldots = r_t^* + rp_t^* - \lambda g_t^*, \tag{12}$$

[12] The ey_t^* decomposition in Result 1 does not separate between discount rates and growth rates in the same manner as a Campbell–Shiller decomposition. Instead, it collects terms such that g_t^* term, for example, contains both the direct cash-flow effect (λg_t^*) and the discount-rate effect (γg_t^*). We do so given our desire to decompose risk-free discount rates into underlying structural components rather than composite terms.

assuming the individual limiting values exist as $t + \tau \to \infty$. In addition, r_t^* satisfies the same decomposition as in (8). Result 1 therefore holds exactly, and the same takeaways apply.

These steady-state facts can be equivalently stated as applying to one-period-ahead conditional expectations:

$$ey_t = \rho_t + (\gamma - \lambda)\mathbb{E}_t[g_{t+1}] + (rp_t - L_t(M_{t+1})),$$

and similarly for the risk-free rate as in equation (7). These versions are useful for interpreting higher-frequency changes in rates and prices.

Case III (Stationarity): Finally, we assume that ey_t and all fundamental variables are stationary, with no unanticipated permanent shocks. This case does not admit permanent changes to real rates or valuations. So to non-trivially characterize the pass-through from interest-rate changes to equity valuations, we must reinterpret all the previous starred terms so as to measure persistent rather than permanent variation. Concretely, we redefine the starred terms $z_t \in \{r_{t+1}^f, rp_t, \rho_t, g_t, L_t(\cdot)\}$ as Campbell–Shiller-type discounted sums:

$$z_t^* \equiv (1 - \delta) \sum_{\tau=0}^{\infty} \delta^\tau \mathbb{E}_t[z_{t+\tau+1}], \tag{13}$$

where $\delta \in (0, 1)$ is a loglinearization term defined in Appendix A.3. Since $(1-\delta)\sum_{\tau=0}^{\infty} \delta^\tau = 1$, (13) defines the starred long-run terms as weighted averages of all future expected outcomes.

Given (13), we again follow Gao and Martin (2021) to obtain the loglinear approximation

$$ey_t^* \equiv ey_t = r_t^* + rp_t^* - \lambda g_t^*,$$

exactly as in (10), where $r_t^* = \rho_t^* + \gamma g_t^* - L_{t,M}^*$ and $rp_t^* = L_{t,M}^* - L_{t,MR}^*$. So given the redefinitions in (13), Result 1 again applies exactly as stated.

In all three cases, therefore, our decomposition generates effectively equivalent results, summarized in Result 1. Only shocks to the pure discounting component of real rates passes through perfectly to equity valuations (in the form of equity yields). Shocks to the other two components in our decomposition — growth rates and uncertainty — generate ambiguous and possibly negative comovement between rates and equity valuations.

2.4 Implications for Equity Duration

Having analyzed the relation between interest rates and equity yields using our decomposition for rates, we now consider what the composition implies for equity duration. Equity duration

does not, as we will see, correspond to the price sensitivity of equity to an arbitrary change in interest rates. Instead, the only interest-rate change that leads to an equity price change equal to its cash-flow duration is a pure discounting shock.

To make this point, we first express equity prices in levels. We consider the constant-growth steady state from Case I and drop time subscripts to simplify:

$$\left(\frac{P}{D}\right)^* = \frac{1}{\exp(r^* + rp^* - \lambda g^*) - 1} = \frac{1}{\exp(\mu^* - \lambda g^*) - 1} \approx \frac{1}{\mu^* - \lambda g^*}. \tag{14}$$

Market-level equity duration $\mathcal{D}$ is defined as the value-weighted time to maturity of the market's expected future cash flows:

$$\mathcal{D} \equiv \sum_{n=1}^{\infty} n \frac{e^{-n(\mu^*)}\mathbb{E}_t[D_{t+n}]}{P} = \frac{1}{1 - e^{-(\mu^* - \lambda g^*)}} \approx \frac{1}{\mu^* - \lambda g^*}. \tag{15}$$

This measure is equivalent to the equity price sensitivity to the log equity discount rate,

$$-\frac{\partial \log P}{\partial \mu^*} = \frac{1}{1 - e^{-(\mu^* - \lambda g^*)}} = \mathcal{D}, \tag{16}$$

which parallels the usual result for the exposure of bond prices to a shift in the yield curve. More important here, though, is price sensitivity to interest-rate changes arising from each of the terms in our decomposition. Given (14), we have the following result describing how price sensitivity depends on the underlying structural driver of interest-rate changes.

Result 2 (Three Interest-Rate Sensitivities)**.** *The sensitivity of stock prices with respect to each of the three terms in the interest-rate decomposition $r^* = \rho^* + \gamma g^* - L_M^*$ is as follows.*

(i) The interest-rate sensitivity of stock prices with respect to pure discount-rate shocks is

$$\mathcal{S}_{r(\rho)} \equiv -\frac{\partial \log P}{\partial \rho^*} = \frac{1}{1 - e^{-(\mu^* - \lambda g^*)}} = \mathcal{D}.$$

(ii) The interest-rate sensitivity of stock prices with respect to growth shocks is

$$\mathcal{S}_{r(g)} \equiv -\frac{\partial \log P}{\partial (\gamma g^*)} = \left(1 - \frac{\lambda}{\gamma}\right) \mathcal{D} < \mathcal{D}.$$

(iii) Assuming that $\partial_{rp,L} \equiv \frac{\partial rp^*}{\partial L_M^*}$ is well-defined and positive, the interest-rate sensitivity of stock prices with respect to risk shocks is

$$\mathcal{S}_{r(L)} \equiv -\frac{\partial \log P}{\partial (-L_M^*)} = (1 - \partial_{rp,L})\, \mathcal{D} < \mathcal{D}.$$

Part (i) tells us that price sensitivity to the pure discount rate pins down equity duration in a manner equivalent to price sensitivity to the equity discount rate in (16). The other interest-rate terms do not share this feature: price sensitivity to growth shocks is strictly less than duration, as is price sensitivity to risk shocks under general assumptions. To take a benchmark example, with log utility ($\gamma = 1$) and equity modeled as an unlevered consumption claim ($\lambda = 1$), the price sensitivity of equity to a change in rates due to g^* or L_m^* are both exactly zero. More generally, as long as equities move positively with expected growth and negatively with respect to risk (as is commonly assumed), the interest-rate sensitivities in parts (ii)–(iii) will both be negative. So only a change in rates induced by a shock to ρ^* moves equities in line with their duration, and equity duration is *not* equivalent to price sensitivity to an arbitrary change in r^*.

The fact that equity duration is equal to price sensitivity to pure discount-rate changes also provides a novel avenue for measuring duration on an ex-ante basis. Measuring duration using realized growth rates, as in the definition (15), requires a very long sample for statistical precision and is inherently backward-looking. Measurement using price sensitivity to μ^*, as in (16), is challenging given the difficulty measuring expected equity returns well. So if our interest-rate decomposition generates reliable estimates of the pure discount rate term ρ^* over time, then estimating the loading of equity returns onto changes in this discounting parameter would allow for clean estimation of duration, both for the market as a whole and for individual portfolios. We pursue this approach in our empirical estimation, which we turn to now.

3. Empirical Implementation

We now implement our real-rate decomposition empirically and study how interest rates and their three components transmit to equities. We do so in a panel of countries, with our data and measurement approach laid out in Section 3.1. We then estimate the terms in our decomposition, first in levels to study secular trends (Section 3.2), and then in changes to study transmission to equity returns (Section 3.3) and portfolio returns in the cross-section of stocks (Section 3.4).

3.1 Data and Measurement Approach

Recall that our goal is to measure each of the terms in the trend real-rate decomposition from Result 1 (or equation (5)): $r_t^* = \rho_t^* + \gamma g_t^* - L_{t,M}^*$, where ρ_t^* is the pure discounting term (or rate of time preference), g_t^* is long-term expected output growth, and $L_{t,M}^*$ is uncertainty (entropy).

Our approach will be to measure r_t^* and g_t^* as directly as possible; measure $L_{t,M}^*$ using a proxy from option prices; and then back out the pure discounting term as a residual.

For trend real rates and expected growth rates, our main input is a panel of long-term forecast data obtained from Consensus Economics. Consensus Economics is a private firm that collects and publishes survey expectations of country-level economic and financial variables by professional forecasters. These forecasters include professional economists at large investment banks and firms, with 10–30 forecasters per survey for each country.[13] We use the long-term forecasts, which are available for the G7 countries (Canada, France, Germany, Italy, Japan, the U.K., and the U.S.) from 1990 through 2023, and for a subset of other developed economies (Netherlands, Norway, Spain, Sweden, and Switzerland) starting in 1995 or 1998. These long-term forecasts were conducted twice annually, in April and October, for the years 1990–2013. Since 2014, the forecasts are available quarterly. For all relevant series, we use consensus (mean) forecasts at the five-year horizon.

To estimate the long-term real rate $r_{t,j}^*$ for date t and country j, we take the consensus forecast of the 10-year nominal interest rate at the end of year $t+5$ and subtract the consensus inflation forecast at the same horizon.[14] For the expected growth rate $g_{t,j}^*$, we use the forecast of real output growth for year $t+5$. One possible concern with this approach is in the potential for a mechanical relation between expected growth and real rates, which might arise if forecasters use a model tying these two variables together when producing their forecasts. While such a mechanical relation (in the absence of a true relation) is of course possible, two points are worth noting. First, a sizable share of the forecasters work at large financial institutions with a key role in trading and pricing assets, so their expectations are likely to be relevant for asset prices irrespective of how they are formed. Second, our main exercise will be to use our decomposition to measure transmission to traded equity prices. These prices will thus allow for an out-of-sample validation of our real-rate decomposition, by testing whether the measured $r_{t,j}^*$ components transmit to equity in the manner predicted by our theory.

To proxy for the uncertainty term $L_{t,M,j}^*$, we build on results from Section 2.3 and Martin (2017). As shown in Appendix A.3, if the market is growth-optimal and the distribution of log growth is symmetric, then the entropy of the SDF is equivalent to that of the market return, $L_{t,M,j}^* = L_{t,R,j}^*$. And as shown in Result 3 of Martin (2017), the squared VIX index is proportional to the risk-neutral entropy of the market return. These imply that setting $L_{t,M,j}^* \propto \text{VIX}_{t,j}^2$ is likely to provide a reasonable approximation, and the constant of proportionality will be implicitly estimated in our regressions for real rates below. We will also

[13] For a list of forecasters for a recent U.S. survey, for instance, see https://web.archive.org/web/20250314034328/https://www.consensuseconomics.com/what-are-consensus-forecasts/.

[14] Our estimates roughly match those of Bauer and Rudebusch (2020), as can be seen by comparing Figure 2 with their Figure 2. The main distinction is a difference in levels, as we consider longer-term rates.

use this squared VIX term as our proxy for the risk term in the equity yield decomposition, as would hold, for example, in the growth-optimal case or with lognormality.

To measure $\text{VIX}^2_{t,j}$, we use a global panel of index option prices from OptionMetrics. The sample, data filters, and implementation approach are taken from Gandhi, Gormsen, and Lazarus (2023); see that paper for details. We calculate the squared VIX directly by implementing the VIX formula for the observed option prices. Our version of the VIX is at the six-month horizon. This is longer than the 30-day horizon calculated by the CBOE for the U.S. market, given our desire to estimate longer-horizon uncertainty. Given the lack of liquid longer-term options, though, we cannot calculate something closer to a five-year VIX. Gandhi, Gormsen, and Lazarus (2023) show that implied volatility decays slowly at longer maturities, so we view our six-month proxy as a reasonable starting point for longer-term uncertainty, and our regression will again scale this value as needed to explain its contribution to real rates. We provide measurement details in Appendix B.1.

For equity prices and valuation ratios, we use a value-weighted index for each country using data from CRSP and Compustat (via the XpressFeed global database). For equity yields $ey_{t,j}$, we start with the five-year earnings-to-price ratio $\overline{E}_{t-4,t,j}/P_{t,j} = [(E_{t-4,j} + \ldots + E_{t,j})/5]/P_{t,j}$, where earnings and prices are calculated on a value-weighted basis for all available traded stocks in the country; see Appendix B.1 for details. We map the earnings yield to the equity yield considered in our theory by multiplying this by 0.5, which is the unconditional average payout ratio in our sample; only around 50% of earnings (more precisely, 49.4%) are paid out to shareholders in an average year–country observation, with the remainder reinvested. An alternative to this approach is to use dividend yields directly as our measure of equity yields, which gives almost identical results ($\text{Corr}(\Delta dp_{t,j}, \Delta ey_{t,j}) > 0.8$ for the full sample).

For our cross-sectional analyses, we use returns on duration-sorted portfolios via Gormsen and Lazarus (2023). This paper measures duration based on analyst forecasts of long-term expected earnings growth (via IBES), or LTG (with higher cash-flow growth indicating a longer duration). We also obtain data on value-sorted portfolios via Ken French's website.

3.2 Secular Trends

To study long-term trends, we start by estimating our decomposition for trend real rates in levels. For all available dates t and countries j, we estimate a regression

$$r^*_{t,j} = \rho_0 + \gamma\, g^*_{t,j} + \beta\, \text{VIX}^2_{t,j} + \Gamma_j + \varepsilon_{t,j}, \tag{17}$$

with country fixed effects Γ_j. In our main specification, we also allow the VIX^2 loading β to differ by country (β_j). While this is not important for our main results, it helps account for

Table 1: Regressions for Trend Real Rates $r^*_{t,j}$

	(1) U.S.	(2) All	(3) All
Expected growth $g^*_{t,j}$	1.8***	2.1***	2.1***
	(0.2)	(0.2)	(0.2)
Uncertainty $\text{VIX}^2_{t,j}$	-10.1**	-3.8	β_j
	(4.5)	(3.0)	
Constant	-1.9***	-1.9***	-2.0***
	(0.5)	(0.4)	(0.4)
Country FEs	✗	✓	✓
Country-Specific $\text{VIX}^2_{t,j}$ Loading	✓	✗	✓
Obs.	86	932	932
R^2	0.57	0.65	0.66
Within R^2	—	0.60	0.61

Notes: This table shows estimated OLS coefficients in the regression (17), along with standard errors in parentheses. In column (1), standard errors are obtained using a block bootstrap. In columns (2)–(3), standard errors are clustered by country and date. Statistical significance at the 10% level, 5% level, and 1% level are denoted by *, **, and ***, respectively. In column (3), the country-specific loadings on the squared VIX, β_j, are statistically significant at the 1% level for 9 of the 12 countries in our sample. The sample is 1990–2023, or the longest available span for the given country.

cases in which sovereign credit risk affects a country's r^*_t.[15] It also allows for the possibility of country-specific measurement error in the VIX, which may be an issue particularly for countries with less-liquid option markets.

Given a set of estimated coefficients and OLS residuals, we then back out the implied pure discounting term as

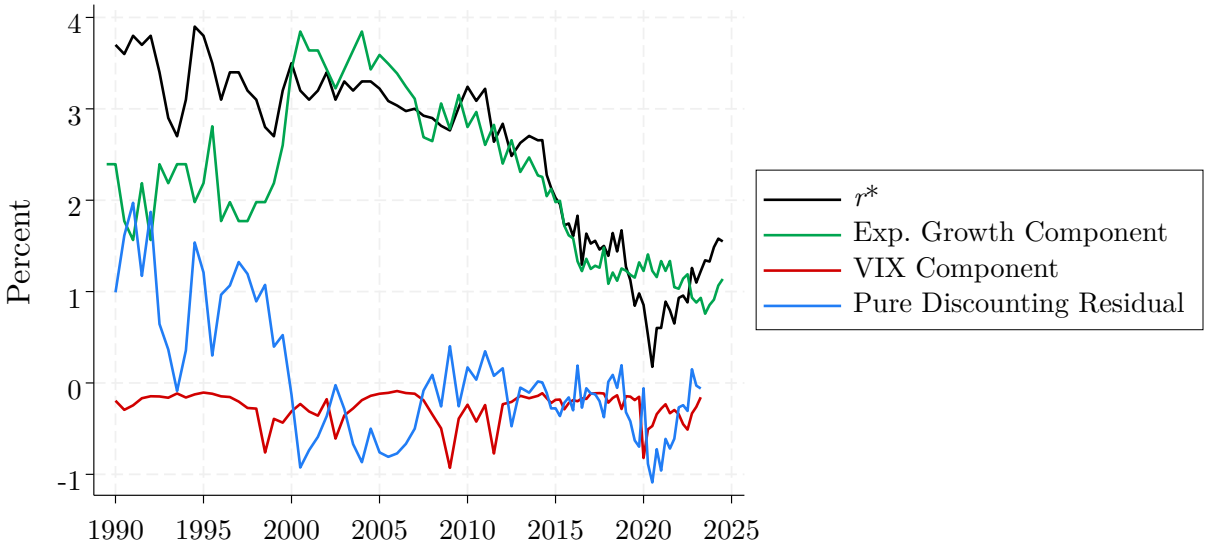
$$\widehat{\rho}^*_{t,j} = \widehat{\rho}_0 + \widehat{\Gamma}_j + \widehat{\varepsilon}_{t,j}. \tag{18}$$

We thus have, by construction, that $r^*_{t,j} = \widehat{\rho}^*_{t,j} + \widehat{\gamma} g^*_{t,j} + \widehat{\beta}\text{VIX}^2_{t,j}$, which corresponds exactly to our theoretical decomposition (with the uncertainty term $-L^*_{t,j}$ proxied by $\widehat{\beta}\text{VIX}^2_{t,j}$).

Estimates for the regression (17) are shown in Table 1, first for the U.S. only and then for the full 12-country panel. The estimates correspond well to our theory. The loading on expected growth is strongly positive and consistently estimated to be close to a value of 2, corresponding to implied relative risk aversion of $\gamma \approx 2$ and intertemporal elasticity of substitution of about 1/2. The loading on the VIX is negative and significant in the U.S. case and for most countries in the country-specific case shown in column (3). This limited set of

[15] Our theory for risk-free real rates suggests that the loading on uncertainty should be negative, but credit risk can induce an offsetting positive relation between risk and long-term rates. We find that this effect is small on average.

Figure 2: U.S. Estimation Results for Decomposition of r^* in Levels



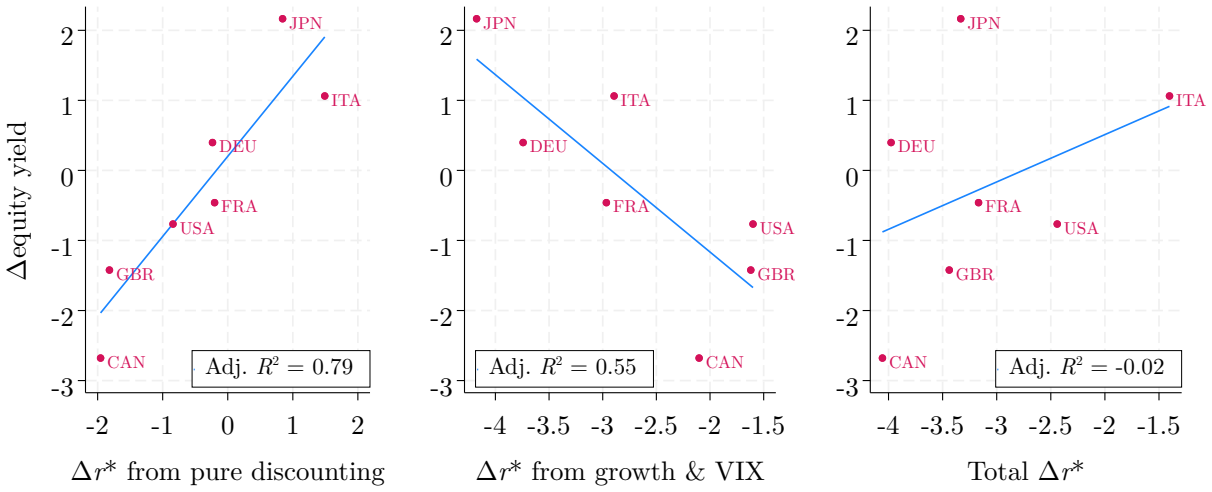
Notes: This figure shows the U.S. trend real rate $r^*_{t,j}$ and its components over time, estimated using (17)–(18) following the main specification in column (3) of Table 1. For readability, the expected growth component is shifted down by 3 percentage points ($\widehat{\gamma} g^*_{t,j} - 3$), and the pure discounting residual is plotted as $\widehat{\varepsilon}_{t,j}$.

variables explains a large share of the variation in trend real rates, with R^2 values of around 0.6 within-country and slightly higher overall. The remaining variation is then attributed to the pure discounting residual.

To visualize the data, Figure 2 presents the estimation results for the decomposition of $r^*_{t,j}$ over time in the U.S. data. The trend real rate has fallen by close to 2.5 percentage points (pp), or 250 basis points (bps), from the beginning to the end of the sample, starting near 4% and ending near 1.5%. As can be seen in the green line, a large share of this decline is attributed to a decline in long-term expected growth. Expected growth fell by around 0.75 pp over the sample, which when multiplied by $\gamma \approx 2$ translates to a predicted decline in yields of about 150 bps. While uncertainty affects real rates during deep recessions, it has little long-term effect over the full sample. The change in growth rates and uncertainty accordingly predicted a decline in real rates of around 150 bps overall, so the additional 100 bps of unexplained decline is attributed to the pure discounting residual. This residual was particularly important in explaining the decline in interest rates early in the sample. From 2000 onward, the decline in interest rates has been driven almost exclusively by declines in expected growth rates, implying little impact on equity valuations.

The remainder of this subsection studies how secular changes in equity valuations across countries relate to changes in the different components of interest rates. Our goal is to understand country-level changes in equity valuations over our sample period.

Figure 3: Main Results: Long-Term Decomposition



Notes: This figure plots the country-level changes in equity yields against changes in different components of interest rates, estimated using (17)–(18) following the main specification in column (3) of Table 1. The leftmost figure plots changes in equity yields against changes in the pure discounting term; the middle figure plots changes in equity yields against changes in the growth and VIX components; the rightmost figure plots changes in equity yields against changes real rates themselves. The sample is 1990–2023, or the longest available span for the given country. For countries for which we can only measure equity yields starting after 1990 (see Appendix B.1), we calculate both Δequity yield and Δr^* over the same window.

In the leftmost panel in Figure 3, we plot the change in equity yields against changes in the pure discounting term in G7 countries. This is the same figure plotted in the right panel of Figure 1 in the introduction. The figure illustrates that the large majority of the changes in equity yields over this sample can be explained by changes in the pure discounting term in interest rates. We emphasize that the pure discounting term is estimated purely from the interest-rate decomposition in (17)–(18), without the use of equity valuations. As a result, there is nothing mechanical about the tight fit in explaining the country-specific change in equity valuations in the last 35 years. This result, along with other complementary evidence presented below for both long and short horizons, thus serves as a strong out-of-sample validation of the estimates from the interest-rate decomposition.

The magnitude of the relation between equity yields and the pure discounting term is almost exactly equal to that predicted by theory. The figure shows that equity yields decrease by one percentage point for every one-percentage-point decrease in the pure discounting term, as in Result 1. And the pure discounting term explains not only relative changes in equity valuations across countries, but also changes in valuations in absolute terms. The intercept for the fit is very close to zero, which means the average earnings yield has moved by as much as the pure discounting term. This finding need not necessarily imply that other factors influencing valuation ratios — such as growth rates and risk premia — have remained

constant, but it does imply that potential movements in growth rates and risk premia have, on net, not played a significant role in changing equity valuations on average over this period. In the later analysis, we in fact find that growth rates have generally gone down and risk premia have gone up.

The middle plot in Figure 3 illustrates the relation between earnings yields and the change in interest rates induced by changes in expected growth rates and uncertainty, taken together. As expected, we find that valuation ratios have dropped in countries where interest rates have dropped because of declines in growth rates and increases in risk: while these changes have decreased interest rates, they have also depressed growth rates on equities and increased equity premia, with the predicted effect on equity valuations being negative. This relationship is noisier than the one plotted in the left panel, consistent with the more ambiguous theoretical predictions for equity valuations given changes in growth rates and uncertainty. But the negative relationship is nonetheless at least moderately strong in the cross-section of G7 countries.

How can the negative relation in the middle panel be squared with the fact that the pure discounting change can nearly perfectly explain the change in equity valuations over time (as documented in the left panel)? Two aspects of the results help in interpreting this. First, note that the best-fit line in the middle panel does not pass through the origin: unlike the Δequity yield—Δpure discounting relationship in the left panel (which features an intercept indistinguishable from zero), the line in the middle panel is shifted by 2.9 percentage points to the left.[16] Enforcing an intercept of zero in this Δequity yield–$\Delta\widehat{r}^*$ relationship, we instead estimate a very small slope (close to -0.1) and an adjusted R^2 of -0.15. Explaining changes in valuations in absolute terms evidently requires using the pure discounting change.

Second, once we account for the pure discounting change in the left panel, the remaining terms in Δr^* do not provide much additional explanatory power for the long-term equity valuation changes. In a regression for Δequity yield on both the pure discounting change and the remaining $\Delta\widehat{r}^*$ terms, only the coefficient on the pure discounting change is significant,[17] and the adjusted R^2 increases only from 0.79 (in the left panel of Figure 3) to 0.81. To visualize this marginal contribution from growth and uncertainty, Figure B.1 in Appendix B.2 shows a version of the middle panel of Figure 3 where the change in the equity yield has now been residualized against the change in the pure discounting term, $\Delta\widehat{\rho}^*_{t,j}$. The part of the equity yield change unexplained by the pure discounting change is generally small quantitatively, and it is now at most very weakly related to the change in rates from the growth and uncertainty

[16] In other words, the average advanced economy had close to no equity valuation change, while nonetheless experiencing growth-rate and uncertainty shocks large enough to decrease real rates by nearly 300 bps.

[17] The estimated loading is 1.7 ($p = 0.048$), while the estimated loading on $\Delta\widehat{r}^*$ is -1.0 ($p = 0.271$).

terms, consistent with the more ambiguous effects predicted theoretically.

Returning to the rightmost panel of Figure 3: while we observe rich comovement between equity yields and the different components of interest rates, we observe almost no relation between equity yields and interest rates themselves. This is because the individual components of interest-rate changes have happened to be somewhat negatively related (albeit weakly so) across countries. As a result, adding the horizontal-axis values in the two left panels of the figure generates a muddled and weak relationship between interest rates and equity yields. This emphasizes how comparing equity valuations to real rates directly can paint a misleading picture.

Discussion and Interpretation

Taken together, Figure 3 provides a clear view of both (i) the secular declines in real rates across countries in recent decades, and (ii) their relation to equity valuations. Taking the U.S. to begin, expected real output growth fell by around 3/4 of a percentage point over the 1990–2023 sample period, and the VIX increased slightly. Given the loadings on these terms in Table 1, those two changes together predict a decline in r^* of about 1.6 percentage points. Instead, r^* fell by 2.5 percentage points. We call the difference of 0.9 percentage points a pure discounting shock, akin to a decrease in the pure rate of time preference. Such a decrease predicts an increase in equity valuations (i.e., a decrease in equity yields), and this is exactly what we see in the left panel of the figure. Taking Japan as a contrasting case, its decline in r^* of 3.3 percentage points is a much smaller decline than would have been expected on the basis of the large decrease in long-term expected growth, indicating a positive pure discounting shock. This positive shock similarly perfectly matches the decrease in Japanese equity valuations. The same applies for all the other countries considered.

While the pure discounting shocks provide a very good description of equity valuation changes in an accounting sense, the question of how to interpret them remains somewhat open thus far. We do not view these changes as likely representing a true aggregate preference (or patience) shock among domestic investors. Instead, a "global imbalances" view of cross-country capital flows, as described by Caballero, Farhi, and Gourinchas (2008), appears to be a reasonable candidate explanation. The main decline in the U.S.'s estimated ρ^* occurred in the mid-to-late 1990s and early 2000s (see Figure 2). This period coincides with a large decrease in the U.S.'s net foreign asset position. Japan's estimated ρ^*, meanwhile, increased during this decade. Strong demand for U.S. assets, particularly from investors in countries experiencing large shocks to the perceived soundness of their financial system (e.g., in the wake of the Japanese stock-market crash), match both the timing and the cross-country patterns observed in Figure 3, as we discuss in greater detail in Section 4.2 below.

3.3 Higher-Frequency Changes and Forecasting Regressions

Interest-rate movements influence not only secular changes in valuation ratios but also higher-frequency fluctuations. In this subsection, we study how stocks move with the different components of interest rates at a higher frequency. The higher-frequency nature of this exercise allows us to conduct our estimation on a within-country basis, in contrast to the cross-country long-difference plots in Figure 3. It also helps avoid potential concerns regarding spurious comovements between slowly moving variables that might arise for the preceding estimation in levels.

When conducting our higher-frequency analysis, we must balance two considerations. First, we wish to explain price and interest-rate variation for reasonably short holding periods. Second, our estimation needs to allow for inertia in forecasters' long-run growth and interest-rate forecasts, which precludes us from considering, for example, monthly returns (since forecasts are collected at most once per quarter). In our baseline analysis in this section, we consider three-year returns and estimate how these move with each of the components of interest rates.[18]

The starting point for this analysis is a regression for changes in trend real rates, analogous to equation (17) but in differences rather than levels:

$$\Delta r^*_{t,j} = \alpha_0 + \gamma\, \Delta g^*_{t,j} + \beta_j\, \Delta \mathrm{VIX}^2_{t,j} + \Gamma_j + \varepsilon_{t,j}, \tag{19}$$

where Δ denotes a three-year change and where the loading on the VIX term is again country-specific.[19] The residual term $\varepsilon_{t,j}$ is now our measure of $\Delta\rho^*_{t,j}$. Next, given this estimated pure discounting change $\widehat{\Delta\rho^*_{t,j}} = \widehat{\varepsilon}_{t,j}$, we regress three-year value-weighted net market returns on that pure discounting term, the change in expected growth, and the change in the VIX, along with a country fixed effect:

$$r^{\mathrm{mkt}}_{t,j} = \alpha_1 + \pi_\rho\, \widehat{\Delta\rho^*_{t,j}} + \pi_g\, \Delta g^*_{t,j} + \pi_V\, \Delta \mathrm{VIX}^2_{t,j} + \Lambda_j + \nu_{t,j}. \tag{20}$$

Table 2 shows the resulting estimates. Before considering our main estimates resulting from (20), we start with a simpler exercise as a benchmark for comparison: we regress three-

[18] In additional analysis, we find that our takeaways are robust to the use of longer or somewhat shorter horizons, though the relationships weaken at horizons shorter than two years (indicating some inertia or measurement error).

[19] Coefficient estimates for regression (19) are presented in Table B.1 of Appendix B.2. The estimates are similar to the level estimates in Table 2 at a high level, albeit with smaller estimated coefficients. This suggests the potential for attenuation bias from measurement error that is amplified when estimating in differences, as highlighted by Griliches and Hausman (1986) and Cochrane (2018). We thank Emi Nakamura for helpful discussions related to this point.

Table 2: Regressions for Three-Year Stock Returns

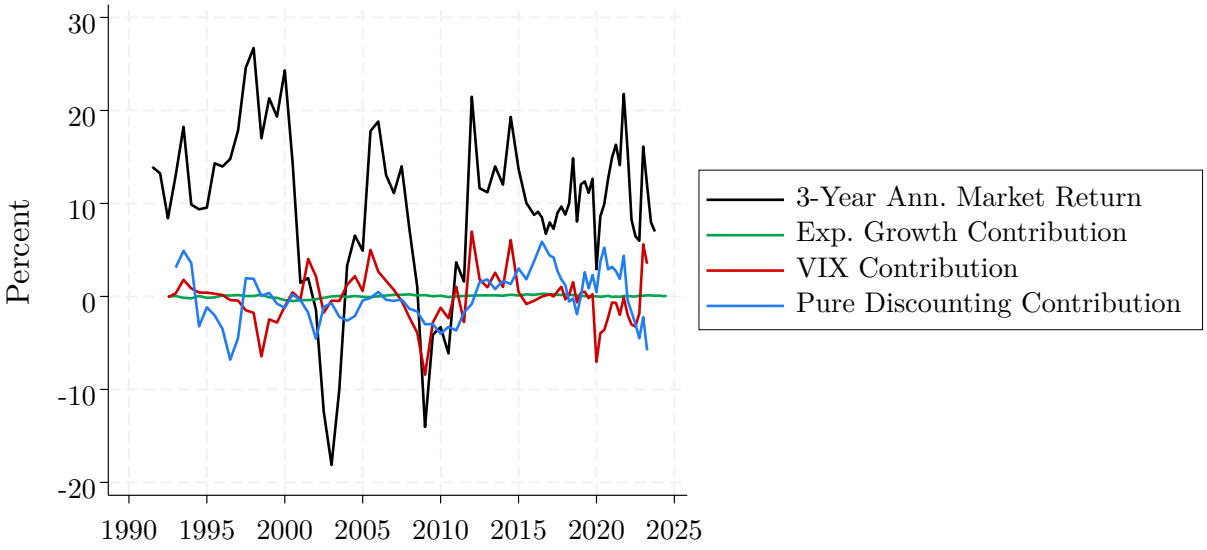
	(1) U.S.	(2) U.S.	(3) All	(4) All
Δ10y yield	4.19		-3.39	
	(3.51)		(2.20)	
Δpure discount ($\widehat{\Delta\rho_t^*}$)		-19.1**		-9.61**
		(7.64)		(3.26)
Δexp. growth		-1.49		16.9*
		(14.0)		(8.82)
$\Delta \text{VIX}^2 \times 100$		-3.08**		-5.44***
		(1.33)		(0.90)
Country FEs	✗	✗	✓	✓
Obs.	74	74	781	781
R^2	0.04	0.20	0.05	0.27
Within R^2	—	—	0.02	0.24

Notes: This table shows estimates from regressing three-year value-weighted market returns on changes in 10-year nominal yields (in columns (1) and (3)), and on changes in the three interest-rate components (in columns (2) and (4)). The interest-rate components are estimated from (19), and the table presents estimates from (20). Columns (1)–(2) consider the U.S. only, while (3)–(4) consider the full panel of developed countries (and include country fixed effects). In columns (1)–(2), standard errors are obtained using a block bootstrap, with block length of one year and 10,000 bootstrap samples. In columns (3)–(4), standard errors are clustered by country and date. Statistical significance at the 10% level, 5% level, and 1% level are denoted by *, **, and ***, respectively. The sample is 1990–2023, or the longest available span for the given country.

year stock returns on the unadjusted change in the traded 10-year nominal yield. Column (1) shows the resulting estimate for the U.S. sample. The slope coefficient is close to zero and statistically insignificant, reflecting the well-known fact that returns on stocks and bonds are close to uncorrelated.

In column (2), we present coefficient estimates from (19) in the U.S. data, showing how stock returns load on each of the three drivers of interest-rate changes: changes in the pure discount term, changes in expected growth, and changes in risk. The loading on the pure discount term is -19 and statistically significant. This loading suggests that stock returns go down by 19 percentage points when trend real rates increase by 1 percentage point due to pure discounting. As in Result 2(i), this coefficient has a clear structural interpretation: it is equal to the negative of the cash-flow duration of the overall market. This market-level duration is often approximated by the dividend yield, generating an estimate on the order of 40 years (see, e.g., Gormsen and Lazarus, 2023). While somewhat lower than that figure, the estimate of 19 years from column (2) is of the same rough order of magnitude and reinforces that the market is a long-duration claim. We view the estimate as a lower bound given

Figure 4: Decomposition of U.S. Value-Weighted Equity Returns



Notes: This figure shows three-year annualized average returns for the value-weighted U.S. stock market, along with estimated underlying contributors. Each contribution term is equal to the estimated coefficient in (20) times the corresponding predictor: for example, the expected growth contribution is equal to the three-year change $\Delta g^*_{t,j}$ times the estimated coefficient $\widehat{\pi}_g$. The coefficient estimates are taken from column (2) of Table 2, but divided by 3 (e.g., the growth loading is -0.5 rather than -1.49). This is to account for the use of annualized returns in this plot, whereas the outcome variable for Table 2 is cumulative non-annualized returns. The pure discounting predictor for (20) is obtained from the first-stage estimation in (19).

the potential for attenuation bias when using the higher-frequency variation in the pure discounting term (see footnote 19 for related discussion).

The remaining estimates in column (2) show that stock returns load very weakly on expected-growth changes, and significantly negatively on changes to risk, again consistent with our theory. Moving to columns (3) and (4) of Table 2, the results are largely similar in the global sample. The main distinction is that the slope on the pure discounting shock is smaller in the global data than in the U.S. data. This lower slope could conceivably reflect measurement issues in the higher frequency data outside the U.S., which attenuates the slope coefficient further.

These higher-frequency estimates allow for a time-series accounting of the period-by-period contribution of different interest-rate components to stock returns. Figure 4 illustrates this higher-frequency return decomposition for the U.S. stock market: it plots the three-year annualized value-weighted market return, along with each of the three fitted components $\widehat{\pi}_\rho \, \widehat{\Delta \rho^*_{t,j}}$, $\widehat{\pi}_g \, \Delta g^*_{t,j}$, and $\widehat{\pi}_V \, \Delta \mathrm{VIX}^2_{t,j}$ (i.e., each predictor variable multiplied by the corresponding loading implied by column (2) of Table 2). Shocks to risk, as shown in the red line, appear more relevant for higher-frequency market returns than was the case in Figure 2 for lower-frequency changes in r^*. This is consistent with the fact that shocks to future equity discount

Table 3: Forecasting Regressions for Future Three-Year Market Returns

	(1)	(2)	(3)
10y yield	0.08		
	(0.38)		
Survey-based r_t^*		0.50	
		(0.68)	
Pure discounting term $\widehat{\rho}_t^*$			2.08***
			(0.61)
Country FEs	✓	✓	✓
Obs.	1,050	842	842
R^2	0.06	0.03	0.06
Within R^2	0.00	0.00	0.03

Notes: This table shows coefficient estimates from forecasting regressions $r_{t,t+3}^{\text{mkt}} = \alpha + \beta X_t + \varepsilon_{t,t+3}$, where $r_{t,t+3}^{\text{mkt}}$ is the country-level annualized three-year market return, and X_t is an ex ante predictor variable. The first column uses the 10-year nominal yield as the predictor variable, using data obtained from each country's central bank. The second column uses our survey-based measure of the trend real rate r_t^* as predictor. The third column uses our estimated pure discounting term $\widehat{\rho}_t^*$, estimated using (17)–(18) following the main specification in column (3) of Table 1. Each regression includes country fixed effects, and all standard errors are clustered by country and date. The sample is 1990–2023, or the longest available span for the given country.

rates and risk premia explain a large share of stock returns (Campbell, 1991), and our use of the VIX as an entropy proxy likely understates the share of the return variation attributable to risk-premium shocks. Expected growth rates, shown in green, do not explain a significant share of the variation in returns in this exercise. The pure discounting contribution, in blue, varies less dramatically than overall returns, but our estimates suggest that it has affected returns significantly in certain subsamples. The increase in the pure discounting term corresponding to the rise in rates since 2021, for example, is estimated to have provided significant headwinds to equities. Returns were nonetheless reasonably high as a result of contemporaneous decreases in equity premia. Such an exercise allows for an assessment of the contributors to stock returns, and their relation to interest rates, on an ongoing basis.

As a final out-of-sample validation test for our interest-rate decomposition in explaining aggregate market returns, we ask whether our estimated pure discounting term $\widehat{\rho}_t^*$ predicts *future* equity returns (in addition to helping account for contemporaneous realized returns). Long-horizon expected equity returns are equal to $\mu_t^* = r_t^* + rp_t^*$. Given that the uncertainty component of r_t^* is likely to be negatively correlated with the equity risk premium rp_t^*, interest rates by themselves are unlikely to be useful for predicting future realized returns. The pure discounting component of r_t^*, by contrast, strips out the uncertainty component of risk-free rates, and therefore should align well with future equity returns.

We conduct such predictability tests in Table 3, which shows coefficients from regressions of annualized market returns over the subsequent three years on ex ante yield-related predictors. Columns (1) and (2) show that neither nominal yields nor our measure of r_t^* help predict equity returns. This provides further evidence that risk premia comove negatively with risk-free yields, as discussed as well by Farhi and Gourio (2018). Meanwhile, as can be seen in column (3), the pure discounting term strongly predicts future returns. While the estimated coefficient of 2.08 is somewhat larger than the theoretical prediction of 1, the estimate is sufficiently noisy that we cannot reject a value of 1 at the 5% level. The upshot of this analysis is similar to our findings above: our interest-rate decomposition succeeds at stripping out shocks to risk-free yields with offsetting effects on equity risk premia (or growth rates), leaving us with a useful measure of the pure discounting component of long-term interest rates.[20] This pure discounting term is accordingly a useful counterfactual long-term risk-free rate to use in calculating a duration-matched equity premium; we will return to this insight in Section 5.1.

3.4 Cross-Sectional Portfolios

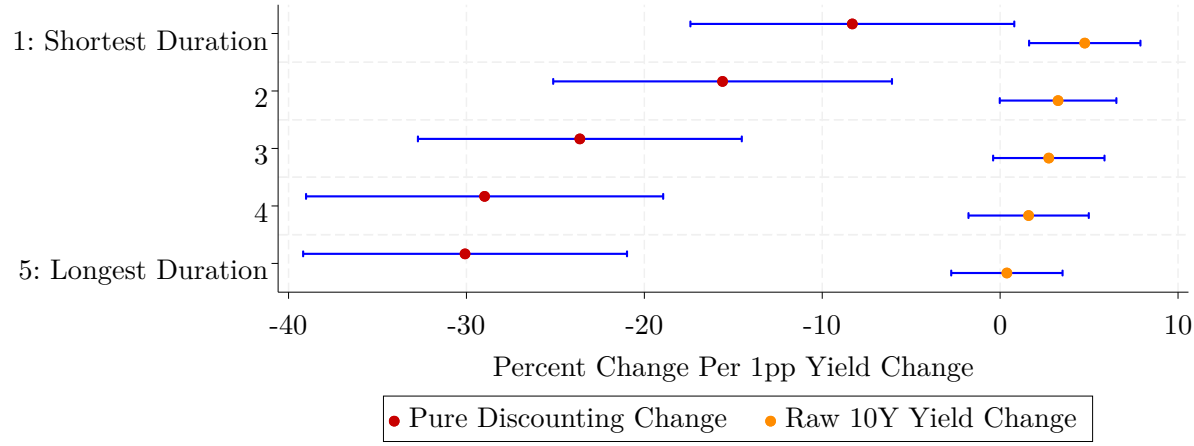
We now turn to the cross-section of stock returns, and study whether firms with different cash flow duration have different exposure to the pure discounting term. A large literature studies the risk and return properties of firms with different cash flow exposure (see Gormsen and Lazarus 2023 and citations therein), finding that firms with shorter cash flow duration have higher risk-adjusted returns. In this section, we use our methodology to quantify cross-sectional differences in cash flow duration, which is a key (and debated) object for this literature.

We focus here on the measure of duration used in Gormsen and Lazarus (2023), which is based on the estimated cash flow growth for different firms. In that paper, we start with analysts' long-term earnings growth (LTG) forecasts obtained from IBES. To extend these forecasts to firms not covered by analysts, we project LTG on a set of contemporaneous firm characteristics; see Gormsen and Lazarus (2023) for details. We then use the fitted values as our measure of predicted duration. According to this exercise, firms with higher cash-flow growth have, all else equal, longer cash-flow duration.[21]

[20] One can also consider a complementary exercise to forecast future *excess* returns. Unlike the version presented in Table 3, such an exercise requires instead stripping out the pure-discounting and expected-growth components of risk-free rates, leaving only a component related to risk (and therefore risk premia). We find in additional tests that this remaining component strongly predicts future excess returns.

[21] Another standard approach is to proxy for duration by valuation ratios (like book to market), since a higher valuation ratio is associated with a longer cash flow duration (i.e., growth firms are long-duration firms). Using book-to-market ratios as the measure of duration does not influence the results presented in

Figure 5: Portfolio Exposures to Pure Discount Rates and Yields: U.S. Stocks



Notes: This figure shows slope estimates from univariate regressions for three-year returns of each of five equity portfolios on the three-year change in (1) the pure discounting term ($\widehat{\Delta\rho_t^*}$), marked in red, and (2) the nominal 10-year yield, marked in orange. Each regression contains a constant. For example, the first point in the top left of the figure shows $\widehat{\beta}_1$ (and 95% confidence intervals) from $r_{t,1} = \alpha_1 + \beta_1 (\widehat{\Delta\rho_t^*}) + \varepsilon_{t,1}$, where $r_{t,1}$ is the three-year return on a value-weighted portfolio of the stocks in the bottom quintile of cash-flow duration. The pure discounting term is estimated from (19) using U.S. data. Duration-sorted portfolios and returns are calculated following Gormsen and Lazarus (2023). The sample is 1990–2023.

In Figure 5, we report slope coefficients of regressions of three-year realized returns onto three-year changes in the pure discounting term for portfolios of U.S. stocks with different cash flow duration.[22] We consider five value-weighted portfolios sorted by duration. The figure shows that the portfolio of firms with the shortest cash flow duration has a slope coefficient of around -10, while the portfolio of firms with the highest cash flow duration has a slope of -30. At face value, these estimates suggest that the cash flow duration of these portfolios varies significantly from -10 to -30 years, which is significant both economically and statistically.

As with the previous analysis, it is possible that the slope coefficients suffer from attenuation bias. If such attenuation bias is driven by classical measurement error, it is similar (in percentage terms) for the different portfolios. In this case, it is useful to focus on the ratio of the cash flow duration of the different portfolios, as this ratio will be unaffected by classical measurement error. We find that the portfolio of firms with longest cash flow duration have three times as long cash flow duration as the firms with the shortest cash flow duration. A lower bound on this difference appears to be 20 years, but we cannot rule out that it is longer.

By contrast, as can be seen in the coefficients plotted in orange, long-duration stocks are not substantially more exposed to raw interest-rate changes than short-duration stocks:

this section.

[22]We present corresponding results for the full global sample in Figure B.2.

all of them have very small estimated loadings when regressing their returns on the change in 10-year nominal yields. And the estimated coefficients go in the "wrong" direction, at least with respect to an interpretation of all interest-rate changes as being exogenous pure-discounting shocks: rather than returns *decreasing* when interest rates increase, they instead weakly increase. This further reinforces the point made in Section 2.4: equity duration does not correspond to the price sensitivity of equity to an arbitrary change in interest rates. Instead, only pure discounting shocks induce interest-rate variation that passes through to equity in proportion to its duration. Duration-sorted portfolios should not, and do not, vary significantly in their exposure to nominal interest rates by themselves; instead, they vary only in their exposure to pure discount-rate changes.

4. Robustness: Changing Profit Shares and Capital Flows

We now consider the robustness and interpretation of our results given two additional channels that are relevant for yields and prices. First, in Section 4.1, we characterize how our decompositions can be generalized when dividend growth is not proportional to output growth as a result of time-varying profit shares of output. Second, in Section 4.2, we consider how our results should be interpreted in the context of globally integrated financial markets.

4.1 Time-Varying Profit Shares

Our main analysis proceeds from the assumption that dividend growth is proportional to output growth; see Section 2.3. While log dividends and consumption should be cointegrated at a sufficiently long horizon, they of course do not comove perfectly at all dates (or for all medium-term forecast horizons). In the realized U.S. data in recent decades, for example, equity cash flows have outpaced GDP and consumption given increases in the corporate profit share of income. Greenwald, Lettau, and Ludvigson (2025) estimate that such unanticipated increases in the profit share accounted for close to 40% of realized U.S. equity returns since 1989. Given that such a channel may be important in our setting, we consider here how time-varying profit shares affect our analysis.

We first note that our interest-rate decomposition in Result 1, $r_t^* = \rho_t^* + \gamma g_t^* - L_{t,M}^*$, is unchanged, as aggregate growth (rather than equity cash-flow growth) is the relevant outcome for the SDF. For equity cash flows, denote dividend growth by $g_{t+1,d} = d_{t+1} - d_t$. We continue to use g_{t+1} to refer to output and consumption growth. Rather than imposing $g_{t+1,d} = \lambda g_{t+1}$ (where we continue to use g_{t+1} to refer to output and consumption growth), we now allow for an arbitrary dividend growth process. So it may be the case that $\text{Corr}(g_t^*, g_{t,d}^*) < 1$. Our

resulting equity-yield decomposition now becomes

$$\begin{aligned} ey_t^* &= \rho_t^* + \gamma g_t^* - g_{t,d}^* + (rp_t^* - L_{t,M}^*) \\ &= \rho_t^* + \gamma g_t^* - g_{t,d}^* - L_{t,MR}^*. \end{aligned} \tag{21}$$

The implications for a change in each of the terms in r_t^* on equity valuations are effectively unchanged from those in Result 1. Only shocks to the pure discounting term ρ_t^* pass through directly from rates to equity yields. Comovements induced by risk shocks are also unchanged. And though the passthrough of growth-rate shocks may be different than in Result 1, it is still the case that such shocks induce weaker pass-through than pure-discounting shocks as long as $\text{Corr}(g_t^*, g_{t,d}^*) > 0$. While there may now be pure dividend-growth shocks (i.e., changes to $g_{t,d}^*$ without corresponding changes in g_t^*), these are entirely separate from the interest-rate dynamics considered in our empirical decomposition for r_t^*, as we return to shortly.

To reconcile this result with the importance of profit-share shocks as estimated by Greenwald, Lettau, and Ludvigson (2025), note that such changes may increase prices *without* affecting equity yields. Their effect on equity yields depends on changes in the path of expected *future* cash-flow growth. If profit-share shocks have tended to cause unanticipated increases in contemporaneous (i.e., already materialized) cash flows rather than expected future growth rates, then both prices and cash flows will increase concurrently, leaving equity yields unaffected as compared to the baseline analysis.[23]

The other possibility is that equity cash-flow growth expectations have diverged meaningfully from output-growth expectations. As discussed above, we view this as a separate shock to anticipated future profit shares. That is, defining $\pi_t^* \equiv g_{t,d}^* - \lambda g_t^*$ for leverage λ, we can write (21) as

$$ey_t^* = \rho_t^* + (\gamma - \lambda) g_t^* - \pi_t^* - L_{t,MR}^*,$$

and a change in the wedge between dividend-growth expectations and levered output-growth expectations is thus a shock to π_t^*.

While shocks to π_t^* are separate from the interest-rate shocks we study, it is an empirically relevant question whether the dynamics of $g_{t,d}^*$ may have differed from those of g_t^*, and whether forward-looking π_t^* has increased in a manner that has happened to offset declines in expected output growth. We can address this question using two sets of additional forecast data, though in both cases the data are only available for the U.S. rather than the full panel.

[23] As a corollary, however, note that the price exposure to such shocks (corresponding to the price exposures analyzed in Result 2) should be significant.

First, the long-term Consensus Economics forecast data provides nominal corporate profit growth forecasts since 1998 for only the U.S. sample. Using this, we construct a proxy for real $g^*_{t,d}$ by subtracting the inflation forecast from the nominal profit-growth forecast for year $t+5$. In this available post-1998 U.S. sample,[24] profit-growth forecasts have in fact fallen by meaningfully more than output-growth forecasts:

$$\Delta g^*_t = -0.50, \quad \Delta g^*_{t,d} = -1.26.$$

With a leverage parameter of $\lambda \approx 2$ — close to estimated coefficient on expected growth in the real-rate decomposition in Table 1, consistent with the insignificant equity effect estimated in (2) — we accordingly estimate that the change in the expected profit-share term $\Delta\pi^*_t$ has been very close to 0 over this sample. As a result, our main conclusions for the U.S. data are unchanged. Profit-share shocks appear to have materialized mainly as changes in current cash flows rather than expected future growth rates, leaving equity yields close to unaffected. And our estimated effect of the pure discounting residual, and the resulting pass-through of roughly 1/3 of the decline in r^*_t to equity valuations in U.S. data, remains unchanged.

As a secondary check on this analysis, we obtain aggregate long-term earnings growth (LTG) forecasts for U.S. equities from Nagel and Xu (2022). These are analyst forecasts for earnings growth for U.S. equities, which Nagel and Xu aggregate to the index level and convert to real terms by subtracting forecasted inflation. Using this series as our second proxy for $g^*_{t,d}$, for the full sample over which we can implement our real-rate decomposition and measure equity yields, we estimate that $\Delta g^*_t = -0.70$, $\Delta g^*_{t,d} = -0.60$.[25] If one again assumes a leverage parameter of about $\lambda = 2.5$, this implies that there has been a modest increase in the profit-share term $\Delta\pi^*_t$, but our main results are largely unaffected.

While Greenwald, Lettau, and Ludvigson (2025) differ from us in their focus on profit-share effects on prices, our results echo some of their findings. They find that "essentially all of the increase in equity values relative to output from the mid-1990s to the end of the sample" can be accounted for by assuming a fixed ratio of market equity to earnings (p 1093), indicating that shocks to contemporaneous earnings relative to output drive the vast majority of equity price (but not equity yield) movements. This is also reflected in the analysis of Atkeson, Heathcote, and Perri (2024), who find that increases in the ratio of cash flows to

[24] This post-1998 sample corresponds to a period during which the majority of the decline in expected output growth took place in the U.S. forecasts.

[25] We prefer the Consensus Economics data because the LTG-based forecasts are much more volatile than the other forecast series (see Bordalo et al., 2024), so the full-sample differences are highly sensitive to the start and end date. Starting the sample in 1992 instead of 1990 gives $\Delta g^*_{t,d} = -1.88$; meanwhile, ending the sample a year later gives $\Delta g^*_{t,d} = 0.14$. Given such large cyclical LTG forecast variation ($\sigma_{\text{LTG}} = \text{SD(LTG)} = 2.0$, vs. $\sigma_{\text{profit}} = 0.8$, $\sigma_{\text{GDP}} = 0.4$), we view the baseline forecasts as better capturing low-frequency variation.

value added have affected valuation much more than the ratio of price to cash flows. While we do not have direct evidence on the importance of profit-share changes outside the U.S., Atkeson, Heathcote, and Perri (2024) also estimate (see their Figure 7) that the increase in cash flows to value added appears fairly specific to the U.S. data.

Finally, we also rerun our higher-frequency return regressions in Table 2, but now with changes in the two $g^*_{t,d}$ proxies (profit growth and LTG) as explanatory variables in addition to the change in expected output growth. Results are presented in Table B.2 in Appendix B.2. Results are very consistent with those presented in Table 2, with the coefficient on $(\widehat{\Delta\rho^*_t})$ very close to the original estimate of -19.1. The loading on ΔLTG is also significant and positive. Overall, our main results continue to apply even when considering changes in the profit share.

4.2 Cross-Country Capital Flows

While the framework introduced in Section 2 does not explicitly rule out cross-border investment, that analysis proceeded from a country-specific domestic SDF and thus did not characterize how foreign capital flows may affect our analysis. Here, we discuss how such flows affect our decompositions both theoretically and empirically, and we show that capital flows help account for the variation in country-level pure discount rates observed in the data.

Consider two large countries i and j. We refer to i as the domestic (or dollar) economy, and j as foreign. Denote the exchange rate, expressed as dollars per unit of foreign currency, by Q_t. (While we refer to Q_t in currency terms, it should be understood as a real exchange rate.) Denote its log by q_t. Also denote the log forward exchange rate by f_t, where the forward horizon matches that of the real rate.

Define $\tilde{\rho}^*_{t,c}$ as the subjective rate of time preference for the average (or marginal) household in country $c \in \{i, j\}$. We will analyze a shock occurring at date t. Assume that prior to the shock, both countries are in steady state, and their decompositions (8) hold with $\rho^*_{t,c} = \tilde{\rho}^*_{t,c}$,

$$r^*_{t,c} = \tilde{\rho}^*_{t,c} + \gamma g^*_{t,c} - L^*_{t,M,c}, \quad c \in \{i, j\}, \tag{22}$$

so that the pure discounting term is equal to the true rate of time preference in each country.

We then consider a shock $dr^*_{t,j} < 0$ in country j. Assume for simplicity that the shock is permanent, so $dr^*_{t,j} = dr^f_{t+1,j}$. The shock to interest rates could come from any of the three terms in the decomposition (22), or a combination thereof. For concreteness, one can consider the experience of Japan in the 1990s, which — in the wake of its financial crisis — featured a large decline in expected growth alongside what Caballero, Farhi, and Gourinchas (2008, p. 367) describe as a "realization that local financial instruments are less sound than they were once perceived to be." This can be represented as a negative growth shock $dg^*_{t,j} < 0$ combined

with a positive pure discounting shock $d\rho^*_{t,j} = d\tilde{\rho}^*_{t,j} > 0$, such that $d\rho^*_{t,j} + \gamma dg^*_{t,j} < 0$. More generally, any $dr^*_{t,j} < 0$ arising from a shock to country j's domestic variables is permissible.

Under no arbitrage, covered interest parity dictates that $fp_t = r^*_{t,i} - r^*_{t,j}$, where $fp_t \equiv f_t - q_t$ denotes the log forward currency premium. As a result, we have

$$dr^*_{t,j} = dr^*_{t,i} - dfp_t.$$

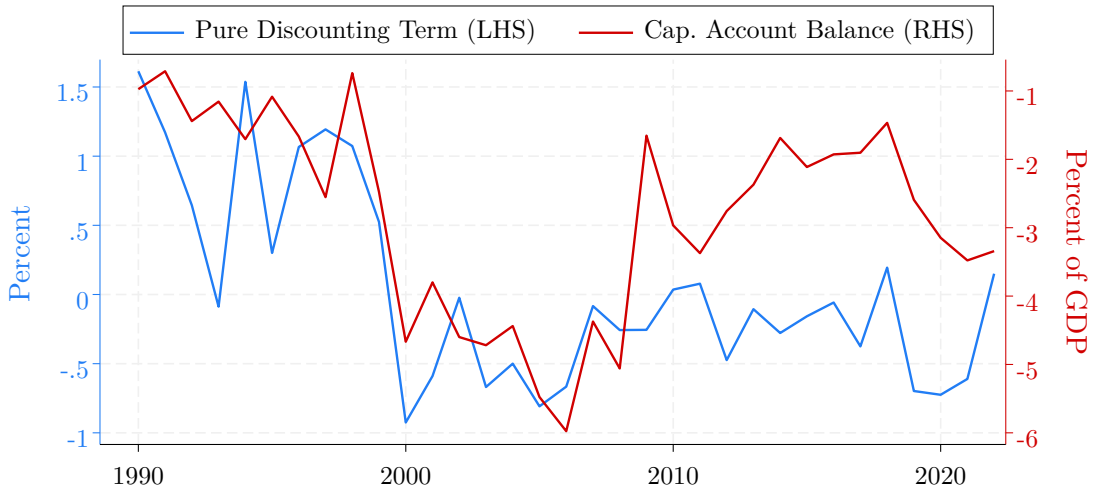
So a decrease in country j's real rate implies some combination of (i) a decrease in country i's real rate, and (ii) an increase in the currency premium for country j relative to i (likely achieved via an immediate depreciation in j's currency and a resulting future forward-implied appreciation).[26] In the absence of a full general equilibrium model (e.g., Alvarez, Atkeson, and Kehoe, 2009, Itskhoki and Mukhin, 2021, or the literature discussed at the end of Section 1), one cannot fully characterize the split between these two channels. But even without fully specifying a supply side, frictions in intermediation or adjustment, or market clearing, one can nonetheless make robust predictions for the effects of $dr^*_{t,j}$ on country i's decompositions.

Consider two extreme cases. First, consider the case in which $dr^*_{t,j} = dr^*_{t,i}$, so the decrease in country j's real rate passes through fully to country i with no change in the forward premium. This would hold, for example, in the case of a currency peg, or approximately in a case in which j is much larger than i or currency intermediation is relatively frictionless. Since $dr^*_{t,i} < 0$ and there are (by assumption) no contemporaneous shocks affecting country i directly, it must be the case on impact that $d\rho^*_{t,i} < 0$: a decrease in country i's interest rate without any corresponding change in forward-looking i-specific expectations must decrease the effective pure discount rate implied by country i's interest rates.

Note that given such a shock, however, $\rho^*_{t,i}$ is no longer equal to the subjective rate of time preference for country i households. Assume that since the shock is permanent, $dL^*_{t,M,i} = 0$. Country i's households' Euler equations still hold, so it must be that $dg^*_{t,i} < 0$ — households immediately consume more given capital inflows (implying lower growth), so that their Euler equations hold — so $dr^*_{t,j} < 0$ leads to both $d\rho^*_{t,i} < 0$ and $dg^*_{t,i} < 0$. That is, a shock that lowers country j's interest rates leads in this case to a lower pure discount rate and lower expected consumption growth in country i. And by the same derivations as in Section 2, the $d\rho^*_{t,i} < 0$ effect will lead to a decrease in equity yields (an increase in valuations), while $dg^*_{t,i} < 0$ will have an ambiguous effect. As long as the survey data effectively measure the change in expected growth in country i, our analysis carries through, and the capital flows from j to i will have the effect of lowering the measured pure discount term in i.

[26] One of these two signs could in principle be flipped with a large enough change in the other term, but all benchmark equilibrium models predict that both terms will be (weakly) of the sign indicated in the text.

Figure 6: Pure Discount Rates and Net Capital Flows in the U.S.



Notes: The line plotted on the left axis shows the end-of-year estimated pure discounting residual $\widehat{\rho}^*_{t,j}$ in the U.S. data, as plotted in Figure 2. The line plotted on the right axis shows the U.S. net capital account balance as a share of GDP, as reported by the IMF (we use their "Net Financial Account" series).

The second extreme case features $dr^*_{t,j} = -dfp_t$, so that the decrease in country j's real rate affects only the forward exchange premium. This would hold if, for example, i is much larger than j, or if currency intermediation is highly frictional so that exchange rates adjust to absorb all demand for country i assets. This case features no change in any of country i's decomposition terms.

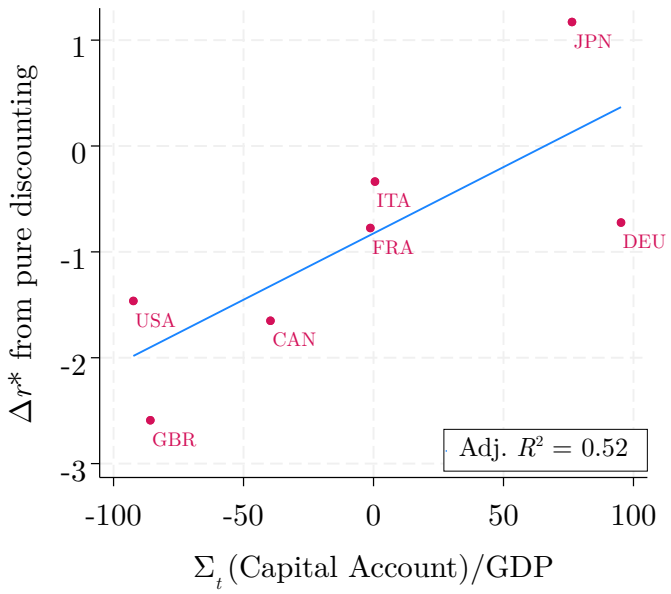
So in either case, as well as in the more empirically relevant intermediate cases, our decompositions still apply in this more realistic setting with international financial markets. A negative shock to a foreign country j's real rate will (weakly) decrease the pure discount rate and the expected growth rate in the domestic country i. As a result, such shocks — and resulting capital flows from foreign to domestic — serve as one (though certainly not the only) potential channel for decreases in the pure discounting term in i.

Motivated by this theory, we examine whether capital flows help explain the pure discount-rate changes in the data. We begin with time-series evidence in the U.S. data. As Figure 6 shows, the variation in estimated ρ^* (plotted on the left axis) aligns well with annual net capital flows (plotted on the right axis, where negative numbers indicate net inflows into U.S. assets).[27] The main decline in the U.S.'s ρ^* occurred in the mid-to-late 1990s and early 2000s. This coincides with a sharp decrease in net capital flows, in line with the "global imbalances" period highlighted and analyzed by Caballero, Farhi, and Gourinchas (2008).

In Figure 7, we analyze whether capital flows help account for pure discount-rate changes

[27]The time-series regression coefficient of $\widehat{\rho}^*$ on capital flows is significant at the 1% level using the robust test recommended by Lazarus et al. (2018). Note that the figure plots $\widehat{\rho}^*$ against annual flows, as Proposition 1 of Caballero, Farhi, and Gourinchas (2008) shows that permanent shocks lead to persistent annual flow effects.

Figure 7: Cross-Country Pure Discounting Changes vs. Cumulated Capital Flows



Notes: The vertical axis shows the country-level change in the pure discounting residual, as plotted on the horizontal axis in the first panel of Figure 3. The horizontal axis shows the cumulative sum of that country's net capital account balance as a share of annual GDP, as reported by the IMF (we use their "Net Financial Account" series). The sample is 1990–2023, or the longest available span for the given country.

in the full panel of G7 countries. There is a fairly strong, positive relation between the full-sample change in a country's estimated ρ^* and the change in its net foreign asset position. As this change in the net foreign position is calculated from cumulated net capital flows, it does not include mechanical valuation effects. That said, capital flows are themselves determined in equilibrium, and they likely have further spillover effects that we do not consider in this illustrative analysis. So while these flows help account for the observed changes in our country-level residuals, we leave a deeper analysis of underlying causal drivers to future work.

5. Additional Empirical Implications

This section uses our framework to shed light on three questions that have been debated in recent literature. We use our methodology to (1) shed light on the "duration-matched" equity premium (Section 5.1), (2) quantify the effect of decreasing interest rates on the value premium (Section 5.2), and (3) help understand the role of an information effect in explaining stock-price responses to monetary policy news (Section 5.3).

5.1 A Significant Duration-Matched Equity Premium

As discussed in the introduction, van Binsbergen (2024) shows that long-term bond portfolios have performed nearly as well as equities in recent decades (see also Andrews and Gonçalves, 2020). In particular, average monthly holding period returns on long-term nominal bond portfolios, constructed to approximate the cash-flow duration of the stock market, have been very close to the average returns on the market since the mid-1990s. So while the premium on the market relative to the short-term risk-free rate has remained high, it appears as if there has been little to no "duration-matched" premium.[28]

The interpretation of this result, however, is less clear. Measuring a duration-matched equity premium is certainly a useful exercise, as it helps in understanding the extent to which high stock returns may have arisen as an essentially mechanical result of the decline in interest rates. But as analyzed in Section 2, bond returns may be high as a result of multiple possible structural drivers, each of which should pass through differently to equities. As a result, an unadjusted nominal Treasury portfolio may not represent an ideal counterfactual long-term bond return for comparison with equities.[29]

van Binsbergen (2024) notes this issue in detail when discussing his results: "the fact that investors have not received compensation for long duration dividend risk does not necessarily mean that investors were not expecting to receive at least some compensation ex ante." In particular, "The results that stocks had poor long-term performance compared to their fixed income counterparts could be driven by a secular decline in long-term real and nominal expected economic growth rates (and/or secular increase in long-term risk premia) over these decades...[but] long-run expected growth measures are not in ample supply." Our data and approach help disentangle the extent to which bond returns should have passed through to equities, and thus how puzzling the apparent low duration-matched equity premium is.

To answer these questions, we construct an alternative counterfactual long-term safe asset return to compare to equity returns. Following the logic above, when seeking to estimate the direct contribution of declining interest rates to realized equity returns, our relevant long-term counterfactual is a duration-matched "pure discounting claim." To understand such a claim, we start by considering a standard zero-coupon bond with maturity n and log yield $y_{t,t+n}$. Its log return from t to $t+1$ can be expressed as

$$r_{t+1,n} = y_{t,t+n} - (n-1)(y_{t+1,t+n} - y_{t,t+n}).$$

[28] As van Binsbergen (2024) states in his conclusion, "One could argue that this simply means that the equity premium puzzle has resolved itself."

[29] Of course, the excess return on equity relative to long-term bonds corresponds to the return from a feasible long-short portfolio; the question is in how to interpret this portfolio's returns.

As usual, the return depends on the initial yield minus the maturity-scaled yield change. A pure discounting claim, by analogy, has return

$$r_{t+1,\rho} = \alpha_t - (n-1)(\rho_{t+1} - \rho_t). \tag{23}$$

The last term in parentheses is the most important for this exercise: our pure discounting claim is constructed so that it appreciates when the pure discounting component of interest rates ρ_t — or, in practice, our estimate of the trend component $\widehat{\rho}_t^*$ — decreases. Because such a decrease should pass through to equity returns in proportion to equity duration $\mathcal{D}$ from Result 2(i), we set the maturity-scaling term to be $n - 1 = \mathcal{D} = 19.1$ years, where the estimate $\mathcal{D} = 19.1$ comes from column (2) of Table 2. As a result, this provides a relevant counterfactual (and in this case, fictitious) long-term bond return to compare to equity returns. The last question is how to define the upfront yield α_t, which determines the level of returns when $\Delta\rho_t = \rho_{t+1} - \rho_t = 0$. Our approach is to set this value to

$$\alpha_t = \widehat{\rho}_t^* + \mathbb{E}_t[\pi_{t+5}], \tag{24}$$

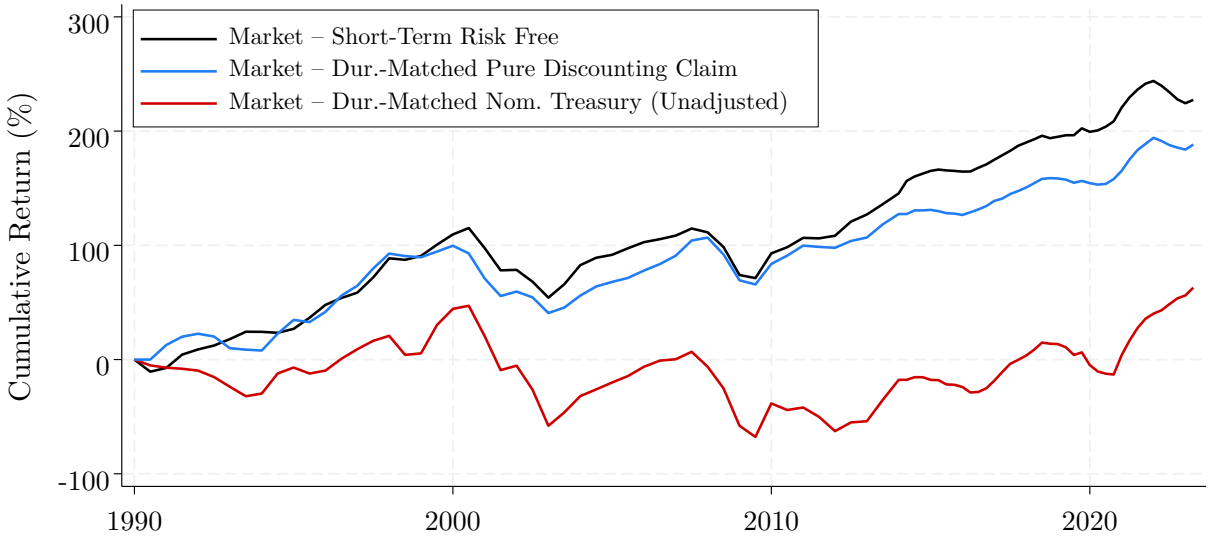
where $\widehat{\rho}_t^*$ is the pure discounting term estimated from (18), and $\mathbb{E}_t[\pi_{t+5}]$ is the consensus survey expectation of long-term inflation (i.e., annual inflation in year $t + 5$). The pure discounting term is defined in real terms, so we add back inflation to put the initial yield in nominal terms (for comparison with a nominal stock return).

We then calculate the excess return on the market relative to this duration-matched pure discounting claim return $r_{t+1,\rho}$, and cumulate returns over time. We compare this duration-adjusted excess return to both (1) the market return in excess of the short-term nominal risk-free rate, and (2) the market return in excess of a duration-matched nominal Treasury security (unadjusted for changes in growth rates or risk), analogous to van Binsbergen (2024).[30]

These three cumulative returns are plotted in Figure 8. As can be seen in the black line, the market has had high average returns relative to the short-term risk-free rate over this period, with a realized annual equity premium of 7.1%. The red line shows a version of the finding in van Binsbergen (2024): when compared to the holding-period returns on long-term nominal Treasuries, much of this premium disappears. The full-sample average return on

[30]Our nominal bond return calculation is somewhat less sophisticated than his. He constructs a bond portfolio with multiple nominal bonds, each weighted in proportion to the value weight of the market's expected future dividend at the corresponding maturity. In contrast, our counterfactual nominal bond return is equal to $r_{t+1,n} = y_{t,t+10} - \mathcal{D}(y_{t+1,t+11} - y_{t,t+10})$. That is, we assume a parallel shift in the yield curve equal to the change in the 10-year nominal yield, and then we calculate the return on a $\mathcal{D} = 19.1$-year Treasury that would result from such a shift.

Figure 8: Cumulative Excess Returns for the U.S. Market



Notes: This figure shows cumulative returns on the value-weighted U.S. stock market in excess of three different counterfactual bond returns. The black line shows the return relative to the short-term risk-free rate. The blue line shows the return relative to the duration-matched pure discounting claim, calculated as in (23)–(24) using the estimated $\widehat{\rho}_t^*$ from (18). The red line shows the return relative to an unadjusted nominal Treasury security with duration $\mathcal{D} = 19.1$ years; see footnote 30 for details of construction.

this nominal-Treasury-adjusted basis is 3.6%. But before considering the last three years of the sample (which featured increasing interest rates and high equity returns), there was no excess return relative to the duration-matched nominal Treasury: the red line in Figure 8 crosses zero in the first quarter of 2021, indicating precisely zero average excess return in the preceding thirty years of the sample.

By contrast, the return on equity in excess of the duration-matched pure discounting claim, shown in blue, is high and stable. On an annualized basis, this realized excess return is estimated to be 6.1% over this period, only slightly (and insignificantly) lower than the standard notion of the equity premium in excess of the short-term risk-free rate. As a result, we estimate a significant duration-matched equity premium once we construct a relevant counterfactual corresponding to the return on a long-term risk-free claim whose appreciation should pass through to equities. We find that it does indeed pass through in the theoretically predicted manner. The return on equity relative to long-term Treasury securities is thus rationalizable: Treasury returns were high in large part because expected growth rates decreased (and, to a lesser extent, because uncertainty increased). The relative performance of stocks over this period is accordingly no longer puzzling once this effect on bond returns is properly accounted for.

5.2 The Value Premium and Interest Rates

We next turn to a puzzling pattern observed in the cross-section of stocks in recent decades. The value premium — measured as the average return on stocks with high book-to-market ratios minus stocks with low book-to-market ratios, or HML (Fama and French, 1993) — has been substantially weaker in recent decades than implied by historical averages. One potential explanation for this underperformance could be that interest rates have dropped, which has led to an unexpected capital gain for the long-duration growth firms, leading growth firms to have performed better than expected ex ante.[31] On the surface, this effect could be meaningful. Imagine that growth firms have a 30-year longer duration than value firms. A naive calculation would imply that a roughly 3 percentage point drop in interest rates would have led to a 90 percentage point relative outperformance of growth firms. Over a 20-year span, this translates to a relative outperformance of more than 4 percent per year, which is large enough to wipe out effectively the entirety of the historical value premium as measured by Fama and French (1993).

The above calculations are, however, not the full story, as discussed in previous sections. First, while interest rates have dropped by close to 3 percentage points in the U.S., the pure discounting term has dropped by only about 1 percentage point, and it is only this component that should pass through to long-duration assets. Second, we estimate that the spread in duration for value-sorted portfolios is substantially below the 30 years assumed above. The net effect on the realized return on the value factor is therefore substantially smaller.[32]

We illustrate and quantify the effect of changes in the pure discounting term for the value factor in the U.S. data in Figure 9. The figure shows both cumulative returns for the HML factor (in black, corresponding to the left axis), and the estimated contribution of changes in the pure discounting component of real rates (in blue, right axis). This pure-discounting contribution is estimated by regressing HML returns on the change in the residual $\widehat{\rho}^*_{t,j}$ from equations (17)–(18), controlling for changes in growth rates and uncertainty,[33] and then multiplying the estimated coefficient by the cumulative change in $\widehat{\rho}^*_{t,j}$ since 1990.

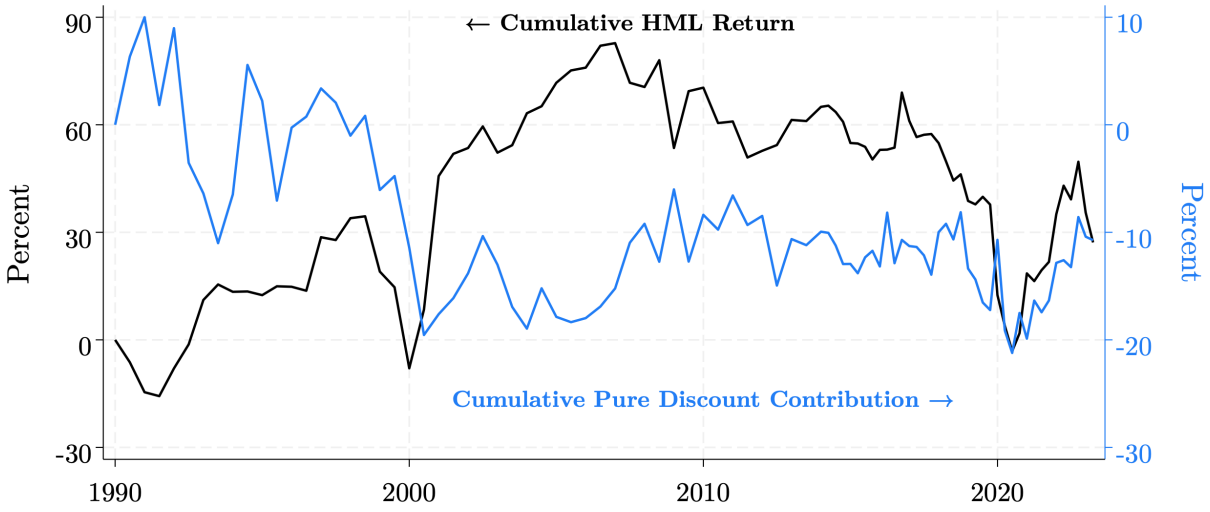
As the figure shows, the effect of the pure-discounting term is modest but non-trivial, reaching a cumulative effect of -20% return at the trough in 2020, but only 10% over the full

[31] This hypothesis is discussed by Maloney and Moskowitz (2021) and Asness (2022), among others.

[32] Consistent with this view, Figure 3 of Gormsen and Lazarus (2023) shows that the return and alpha on a short-minus-long-duration strategy has been quite consistently positive over recent decades.

[33] This estimation exercise parallels the one in equation (20) for the overall market, with the exception that here we use the change in the residual $\widehat{\rho}^*_{t,j}$ estimated in levels from (17) (rather than the residual from the first-difference estimation in (19)). We do so because this allows for straightforward estimation of a cumulative effect of the change in $\widehat{\rho}^*_{t,j}$ in levels, as needed for the exercise in Figure 9. Cumulating the effects of three-year changes $\widehat{\Delta\rho^*_{t,j}}$, by contrast, would only allow for measurement of the pure discounting contribution every three years (so the blue line in Figure 9 would only show 11 equally spaced points).

Figure 9: The Contribution of Pure Discounting Changes to Value Factor Returns

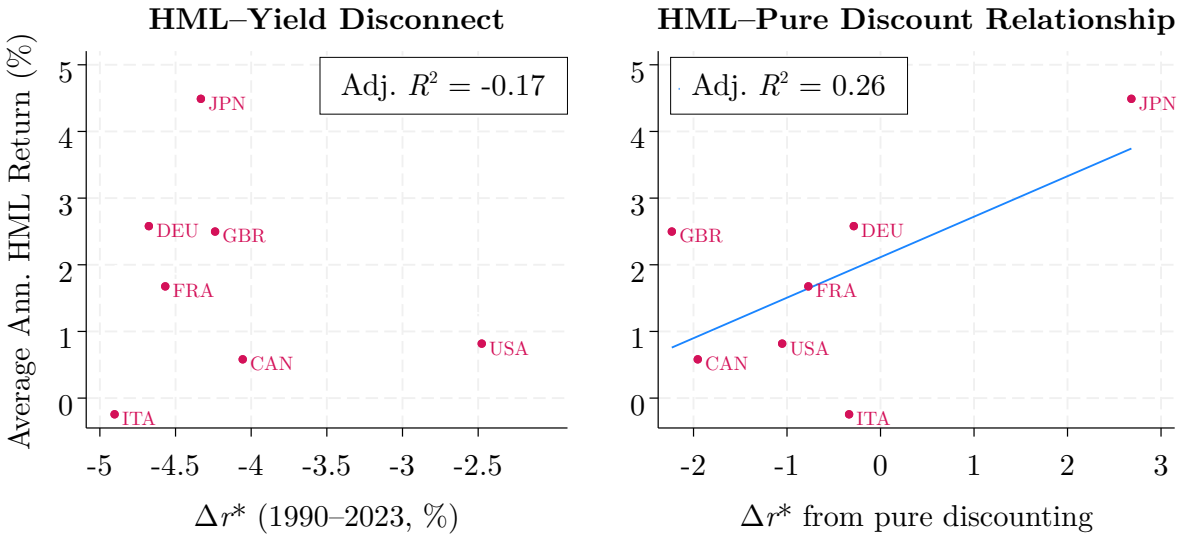


Notes: The black line shows the cumulative return on the Fama and French (1993) HML value factor for the U.S. sample since 1990, obtained via Ken French's website, plotted on the left axis. The blue line shows the contribution we estimate is attributable to the pure discounting component of real rates, plotted in cumulative percent terms on the right axis. For this estimated contribution, we begin with the pure discounting residual $\widehat{\rho}^*_{t,j}$, estimated using (17)–(18) following the main specification in column (3) of Table 1, as plotted in Figure 2. We regress three-year HML returns on the three-year change in this residual, along with the three-year change in expected growth rates and the three-year change in the squared VIX (akin to (20)). The estimated contribution from the three-year residual is then the estimated coefficient on $\Delta\widehat{\rho}^*_{t,j}$ multiplied by the cumulative change in $\widehat{\rho}^*_{t,j}$ since 1990.

sample. In addition, Figure 9 also shows that the crash and rebound of the value factor from 2020–2023 matches the dramatic changes to the pure discount term experienced over those years at least in timing, if not fully in magnitude: the cumulative HML return over that period is around 30% (on the left axis), while the percent attributable to the pure discounting term is around 10% (on the right axis). So while the pure discounting contribution is often important, it is clearly not the full story explaining the performance of value in recent decades in the U.S. sample.

In Figure 10, we exploit our global panel to study what share of the cross-country differences in realized value returns since 1990 can be explained by cross-country differences in the evolution of the pure discounting term. The figure shows that value firms in countries that have experienced a larger decrease in the pure discount term have had lower realized premia relative to growth firms over the sample period. And Japan — which has had an increase in the pure discounting term over our sample — has had the largest realized value premium. The cross-sectional R^2 demonstrates meaningful explanatory power, but it also indicates that the returns to the value factor cannot be fully summarized by changes in the

Figure 10: Interest Rates and Value Returns: Long-Term Global Evidence



Notes: The left panel plots the country-level average annual return on the high-minus-low (HML) book-to-market value factor versus the change in estimated trend real rate r^* over the 1990–2023 sample. Following Fama and French (1993), each country's HML factor is constructed based on a 2×3 size and book-to-market double sort, with returns calculated as the average of the high-minus-low return for small and large firms. See Gormsen and Lazarus (2023) for details. The right panel plots the same average annual HML return against the change in the pure discounting term estimated using using (17)–(18) following the main specification in column (3) of Table 1. The sample is 1990–2023. Note that HML returns are available starting in 1990 for all countries, whereas the equity yield samples start later than this in some cases (see Appendix B.1). As a result, the Δr^* terms in this figure may differ relative to Figure 3, given the earlier start date for the changes calculated in this figure.

pure discounting term.[34] So while this pure discounting change is important for explaining some share of the cross-country value premium, it clearly does not represent the full story over this period.

5.3 Unpacking Monetary Policy Shocks

As a final exercise, we use our decomposition and estimation results to help unpack the effects of surprise changes in short-term interest rates by monetary policymakers. These surprises, when properly measured (e.g., using high-frequency changes in interest rates around policy announcements), are by construction exogenous shocks to short-term nominal rates. These shocks then pass through strongly to long-term nominal and real rates (e.g., Hanson and Stein, 2015). But while some papers have treated the resulting changes in long-term rates as if they represent pure discounting shocks (i.e., shocks to ρ_t), this is not necessarily a valid

[34]We find that the same results hold — both within and across countries — when considering HML alpha (i.e., on a market-adjusted basis), rather than considering the raw value premium.

assumption: while the change in the short-term rate is exogenous, the long-term yield change depends on changes to the pure discount rate *as well as* changes to perceived long-term growth and uncertainty.[35] That is, the long-term real yield response depends not just on the pure short-rate shock (and the perceived persistence of this shock), but also on the market's perceived future changes to endogenous outcomes resulting from this shock.

In practice, if positive interest-rate shocks are contractionary (i.e., decrease expected growth rates) and cause increases in uncertainty, then the observed long-term real yield change $\Delta y_{t,t+n}$ will in fact tend to understate the change in the pure discounting term $\Delta\rho_t$, from (5). In this case, attributing the entirety of the yield change to $\Delta\rho_t$ may be innocuous as a conservative assumption. But in the presence of something like an information effect (Nakamura and Steinsson, 2018), under which positive interest-rate shocks lead the market to revise growth expectations *up*, then the validity of such an assumption is less clear.

A benefit of our framework is that we can estimate directly the perceived effect of any given shock on the separate components of real rates. One approach to this would be to observe the change in expected growth rates around an announcement and then strip out these changes, akin to the approach taken in Section 3. But the timing of the Consensus Economics surveys makes such an approach challenging when considering high-frequency shocks like monetary policy surprises. First, the surveys are conducted infrequently (either every six months or every three months), and forecasters may exhibit inertia in changing their growth-rate forecasts after a given shock.[36] Second, the surveys *prior* to a given monetary policy change may be stale by the time of the FOMC meeting, inducing a possibly spurious positive relation between expected-growth revisions and policy surprises: if positive shocks tend to occur in the wake of good economic news revealed between the most recent Consensus Economics survey and the next FOMC meeting, then this could result in positive revisions that do not reflect the change in forecasts resulting from the announcement itself. See Bauer and Swanson (2023b) for an extensive related discussion.

Instead, we can take advantage of the fact that we observe three high-frequency asset-price changes on the announcement dates themselves: we observe the change in long-term yields

[35]To take one recent example, Kroen et al. (2024) "empirically analyze the impact of falling rates on firms using high frequency interest rate shocks at FOMC announcements as exogenous shifters to the interest rate," and then use this assumption to estimate the duration of market leaders (large stocks) relative to followers (small stocks) based on their relative return responses to interest-rate shocks. Such stock-price changes pin down duration only under the assumption that the shock consists only of a pure discount-rate shock. As we find in the analysis presented here, this assumption is not too far from the truth on average. In addition, in separate analysis, we estimate that large stocks appear to have somewhat longer durations than small stocks (measured using exposure to pure discounting changes) in the recent low-rate sample, which is the relevant subsample for the Kroen et al. (2024) analysis. That said, prior to 2013, small stocks in general have longer estimated durations than large stocks, consistent with the full-sample results in Gormsen and Lazarus (2023).

[36]Partly due to this inertia, our difference-based estimation in Section 3.3 considers three-year changes.

$\Delta y_{t,t+n}$, the return on the market r_t^{mkt}, and the change in uncertainty proxied by ΔVIX_t^2. And our previous estimation provides a mapping from any change in ρ_t, expected growth g_t, and uncertainty VIX_t^2 to a change in long-term yields and stock returns. As a result, this mapping can be inverted to provide an estimate of the change in ρ_t and g_t implied by the observed asset-price changes. For example, a positive market return coinciding with an increase in yields implies that expected growth must have increased by enough, or uncertainty must have decreased by enough, to offset any given increase in the pure discounting term. Given that we can observe the change in uncertainty, these two reactions in fact exactly pin down the required change in both terms.[37]

To implement this idea, we start with a slightly modified version of the yield change decomposition in (19). The benchmark yield change we will use in the high-frequency data is a change in the 10-year yield, whereas (19) was estimated for the five-year yield. We therefore re-estimate that equation using three-year changes in the 10-year trend real yield as our starting point. In practice, the resulting estimates are quite close to those presented in Table B.1.[38] Next, we estimate (20) using the three resulting terms from that decomposition. Estimates are again quite similar to the benchmark shown in Table 2.[39]

We then use data from Bauer and Swanson (2023a), who provide changes in 10-year nominal yields, S&P 500 futures returns, and monetary policy shocks (orthogonalized with respect to ex ante predictors) in 30-minute windows around FOMC announcements.Based on the results of Nakamura and Steinsson (2018), we assume that the change in 10-year nominal yields is equal to the change in 10-year real yields, $\Delta y_{t,t+10}$. Finally, we calculate the daily change in the VIX_t^2 on the announcement day. Using these observed high-frequency changes and our estimated coefficients in the real-rate and stock-return regression, we invert the following two equations for the two unknowns $\Delta\rho_t$ and Δg_t:

$$\begin{aligned}\Delta y_{t,t+10} &= \Delta\rho_t + \widehat{\gamma}\,\Delta g_t + \widehat{\beta}_j\,\Delta \text{VIX}_t^2,\\ r_{t,j}^{\text{mkt}} &= \widehat{\pi}_\rho\,\Delta\rho_t + \widehat{\pi}_g\,\Delta g_t + \widehat{\pi}_V\,\Delta \text{VIX}_t^2.\end{aligned}$$

We then regress the recovered $\Delta\rho_t$ and Δg_t, as well as $\Delta y_{t,t+10}$ and ΔVIX_t^2, on the orthogonalized monetary policy shocks mps_t from Bauer and Swanson (2023a).

Across all post-1994 announcements, a 100-basis-point shock (i.e., a shock scaled so that

[37] Our estimation approach is somewhat similar in spirit, if not in implementation, to the one used by Knox and Vissing-Jorgensen (2024) to decompose contemporaneous changes in observed returns. Our approach requires slightly more structure than the one in Nagel and Xu (2024), but it is again similar in spirit.

[38] The estimated growth-rate loading is indistinguishable from the one presented in column (3) of that table, while the loading on VIX is slightly smaller (-2.3 in the U.S., rather than -4.3).

[39] In this case, they are statistically indistinguishable from the figures presented in column (2), with the very slight change being that the estimated loading on the expected-growth change is now 2.4.

the impact on the one-year Eurodollar futures contract is +100 bps) results in the following:

(i) An increase in the 10-year yield of $\Delta y_{t,t+10} = 45$ basis points (similar to Nakamura and Steinsson, 2018). This result is significant at 1%, and the regression has an R^2 of 0.36.

(ii) An increase in the VIX of $\Delta \text{VIX}_t^2 = 0.013$ (or, in non-squared terms, an increase of 0.2%). This small increase is nonetheless significant at 1%. The regression R^2 is 0.04.

(iii) An increase in the pure discounting term of $\Delta \rho_t = 29$ basis points. This is again significant at 1%, and the regression has a high R^2 of 0.30.

(iv) An increase in expected growth of $\Delta g_t = 7$ basis points. This is significant only at 10%, and the R^2 is 0.04.

As a result, we conclude that the average monetary policy shock indeed appears reasonably close to a pure discounting shock, at least in its effect on long-term yields. This conclusion is fairly similar to that of Nagel and Xu (2024), using different methods. But there is nonetheless a small, somewhat noisily estimated *positive* expected-growth-rate change estimated as resulting from a contractionary shock. Intuitively, while stock returns decrease following contractionary shocks, they do not decrease on average by quite enough — i.e., they decrease by less than 19% (given an estimated duration of around 19 years) for every one-percentage-point change in long-term yields — to be consistent with a pure discounting shock alone.[40] As a result, we find some evidence in favor of an information effect on average, using different methods than those used by Nakamura and Steinsson (2018).

The fact that the R^2 in the growth-rate regression is so low indicates that there is meaningful announcement-specific heterogeneity in the perceived effects on growth rates (as well as the other endogenous variables): there are some announcements with strong conventional policy responses, and others with strong apparent information effect–type responses. For example, the accommodative announcements on March 23, 2020, during the depths of the market downturn at the onset of the Covid crisis, is estimated to have increased long-term expected growth rates significantly. In future work, we plan on unpacking this heterogeneity in greater detail to take advantage of our announcement-specific estimates.

6. Discussion and Conclusion

We provide a new framework and measurement tools to decompose any change in real interest rates into mutually exclusive underlying structural changes. According to our decomposition, only pure discounting shocks should pass through perfectly from real yields to equity valuations

[40] They would in fact need to decrease by more than 19% to be consistent with such a shock, given the small positive effect on the VIX.

theoretically. When implemented empirically with long-term survey forecast data and a panel of asset prices, the decomposition works very well: pure discounting shocks are estimated to pass through one-for-one to equity yields, while the other components of interest-rate changes do not.

The recovered pure discounting component of real rates helps us answer a range of important questions related to asset pricing, macroeconomics, and secular economic trends observed in recent decades. In the U.S. data, we estimate that a sizable share of the decline in interest rates since 1990 — around 35% — is attributable to the pure discounting term, indicating some meaningful pass-through from declining yields to rising risky-asset valuations. But assuming perfect pass-through, as a range of literature has done, nonetheless overstates the effect of declining interest rates by roughly three times. The partial pass-through we find implies that much of the rise in household wealth (and inequality) was likely non-mechanical.[41] Our estimates also imply that stocks have continued to exhibit a sizable equity premium relative to a duration-adjusted counterfactual. In further analysis, we use our decomposition to speak to higher-frequency equity returns, explain interest rates in the cross-section of stocks, and better understand the perceived effects of monetary policy shocks.

Unpacking the drivers of country-level changes in the pure discounting term in a structural sense, over and above the analysis of capital flows in Section 4.2, will be important for better understanding how to interpret these changes. But in spite of the work to be done on this, our paper provides a clear framework and tools to understand the relationship between stocks and bonds. This bond-stock relationship appears chaotic, both at high frequencies and over the long run, as is apparent from the stock–yield disconnect shown in the left panel of Figure 1. But our simple framework, combined with long-term survey data, works very well at isolating a pure discounting component of interest rates that explains both higher-frequency stock returns and longer-term secular changes in equity valuations, as in the right panel of Figure 1.

One implication of our findings is that we can nearly perfectly explain the long-term changes in both interest rates and equities without the need for any additional convenience yield specific to Treasuries. While such market-specific shocks may be quite important for explaining shorter-term fluctuations (as seen, for instance, in Di Tella et al., 2024), "standard" asset pricing evidently works reasonably well at explaining the data at a low frequency.

[41] That said, more work needs to be done to understand the pass-through of interest-rate changes to assets other than equity, which are important for many households' wealth.

Appendix

A. Additional Theoretical Derivations and Discussion

This appendix provides proofs, derivations, and discussion of theoretical results as referenced in the main text.

A.1 Additive Log SDF Decomposition

This subsection discusses the additive decomposition for the log SDF in equation (2), and its relation to the decomposition in Hansen (2012). Given our discrete-time environment, for notational simplicity we will in fact more directly build on Proposition 4.2.1 of Hansen and Sargent (2022). That result provides a discrete-time analogue to the continuous-time decomposition for additive functionals in Theorems 3.1–3.2 of Hansen (2012),[42] which we then apply to our specific setting.

We begin by defining the process S such that $S_0 = 1$ and $M_{t+1} = \frac{S_{t+1}}{S_t}$, so $m_{t+1} = s_{t+1} - s_t$. We consider three separate cases for fundamental dynamics, which parallel the three cases considered in Section 2.3.[43]

A.1.1 Stationarity with Unanticipated Breaks

Following Hansen and Sargent (2022), begin by defining X_t $(t = 0, 1, 2, \ldots)$ to be a stationary Markov process of dimension n with transition equation

$$X_{t+1} = \varphi(X_t, W_{t+1}), \tag{A.1}$$

where $\varphi(\cdot, \cdot)$ is a Borel-measurable function and W_{t+1} is a k-dimensional vector of unanticipated shocks satisfying $\mathbb{E}_t[W_{t+1}] = \mathbb{E}[W_{t+1}|X_t] = 0$. The dynamics in (A.1) induce a transition distribution $\mathcal{P}$ for X.

Given the log SDF $m_{t+1} = s_{t+1} - s_t$, assume that the process s is an additive functional, in the sense that it can be represented as

$$s_{t+1} = s_t + \kappa(X_t, W_{t+1}),$$

where $\kappa : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}$ is a measurable function. Define the unconditional expectation of

[42] See also Hansen (2019).

[43] The first case is a generalization of Case I in Section 2.3, allowing for arbitrary Markov dynamics (rather than the specialized version with conditionally i.i.d. dynamics in Section 2.3) along with unanticipated breaks.

the increment $s_{t+1} - s_t$ to be $\nu = \mathbb{E}[\kappa(X_t, W_{t+1})]$. Define $\overline{\kappa}(x) = \mathbb{E}[\kappa(X_t, W_{t+1})|X_t = x] - \nu$ to be the deviation of the expected increment conditional on $X_t = x$ from its unconditional mean. Using the infinite sum of all such future deviations as of t, define

$$H_t = \underbrace{\kappa(X_{t-1}, W_t) - \nu}_{(s_{t+1}-s_t)-\nu} + \sum_{j=0}^{\infty} \mathbb{E}_t[\overline{\kappa}(X_{t+j})], \tag{A.2}$$

and assume that the sum in (A.2) converges in mean square to a finite variable. Next, define

$$h(X_t) = \mathbb{E}_t[H_{t+1}] = \sum_{j=0}^{\infty} \mathbb{E}_t[\overline{\kappa}(X_{t+j})]. \tag{A.3}$$

Finally, define the martingale increment

$$\varepsilon_{t+1} = H_{t+1} - h(X_t), \tag{A.4}$$

so that $\mathbb{E}_t[\varepsilon_{t+1}] = 0$ by construction.

Given the above setup and the additional assumption that $\mathbb{E}[\kappa(X_t, W_{t+1})^2] < \infty$, Proposition 4.2.1 of Hansen and Sargent (2022) then gives that $m_{t+1} = s_{t+1} - s_t$ satisfies the following additive decomposition:[44]

$$m_{t+1} = \nu + h(X_{t+1}) - h(X_t) + \varepsilon_{t+1}. \tag{A.5}$$

The first term represents the linear trend in s_t. The second component, $h(X_{t+1}) - h(X_t)$, is a stationary difference. The last term is a mean-zero martingale increment.

We now map to our interpretation of the first term as representing discounting and the second term as depending on cash-flow growth. Denote the log cash-flow process by c_t, and assume that

$$\begin{aligned} c_{t+1} - c_t &= \mu_c(X_t) + \sigma_c(X_t) B_c W_{t+1}, \\ X_{t+1} &= A_x X_t + \sigma_x(X_t) B_x W_{t+1}, \end{aligned} \tag{A.6}$$

where $\mu_c(X_t)$, $\sigma_c(X_t)$, and $\sigma_x(X_t)$ are measurable functions. This nests many common specifications for fundamentals. If (i) W_t is a 3×1 i.i.d. standard normal vector, where the first entry contains a shock to contemporaneous cash flows (so only the first entry of

[44] Theorem 3.2 of Hansen (2012) gives an exactly analogous decomposition in continuous time, with the added interpretation of h as a finite-second-moment solution to $\lim_{t \searrow 0} \frac{1}{t}\mathbb{E}_0[h(X_t) - h(x)|X_0 = x] = \overline{\kappa}(x)$, where $\overline{\kappa}(x)$ is the deviation of the local mean of the increment in s from its unconditional mean ν.

the row vector B_c is non-zero), (ii) X_t is a 2×1 vector, where the first entry $X_{1,t}$ affects expected cash-flow growth ($\mu_c(X_t) = \overline{\mu} + X_{1,t}$), the second captures stochastic volatility ($\sigma_c(X_t) = \sigma_x(X_t) = \sqrt{\overline{\sigma}^2 + X_{2,t}}$), and A_x and B_x are 2×2 diagonal matrices, then this represents the long-run risks model of Bansal and Yaron (2004) with stochastic volatility. The Campbell and Cochrane (1999) habit formation model, meanwhile, applies if W_t is scalar i.i.d. standard normal, $\mu_c(X_t) = \overline{\mu}$, $\sigma_c(X_t) = \sigma$, $B_c = 1$, and X_{t+1} is the deviation of the log surplus consumption ratio $\widetilde{s}_t$ from its long-run mean (with $\sigma_x(X_t)$ representing Campbell and Cochrane's "sensitivity function" $\lambda(\widetilde{s}_t)$). Other settings fit similarly within the framework.

To relate the SDF dynamics to cash flows, assume that the log SDF can be written as

$$m_{t+1} = -\rho - \widetilde{\gamma}(c_{t+1} - c_t) + n_{t+1}, \tag{A.7}$$

for some constant $\widetilde{\gamma}$ and process $n_{t+1} = \mu_n(X_t) + \sigma_n(X_t)B_n W_{t+1}$. This is again quite general.[45] As in Section 2.2, it holds in a representative-agent, power-utility setting, where $\rho = -\log\beta$ is the time discount rate, $\widetilde{\gamma} = \gamma$ is relative risk aversion, and $n_{t+1} = 0$. If the representative agent has Epstein-Zin preferences with elasticity of intertemporal substitution (EIS) ψ, time discount rate ρ, and relative risk aversion γ, (A.7) holds, but now with $\widetilde{\gamma} = \frac{1}{\psi}$ and $n_{t+1} = (1/\psi - \gamma)(v_{t+1} - (1-\gamma)^{-1}(\log \mathbb{E}_t[V_{t+1}^{1-\gamma}])$, where V_{t+1} is continuation utility and v_{t+1} is its log.[46] In the Campbell and Cochrane (1999) habit setting, (A.7) holds, with $\widetilde{\gamma} = \gamma$ again representing risk aversion and $n_{t+1} = -\gamma(\widetilde{s}_{t+1} - \widetilde{s}_t)$, where $\widetilde{s}_t$ is the log surplus consumption ratio. It can also be mapped straightforwardly to various heterogeneous-agent models, such as that of Constantinides and Duffie (1996) or subsequent models.

To relate the above representation to the additive decomposition (A.5), we can construct each of the terms in that decomposition under our assumptions on cash flows in (A.6) and the SDF in (A.7). Define $\nu_c = \mathbb{E}[c_{t+1} - c_t] = \mathbb{E}[\mu_c(X_t)]$ and $\nu_n = \mathbb{E}[n_{t+1}] = \mathbb{E}[\mu_n(X_t)]$, and assume that $\nu_n = 0$. This will hold as long as the additional perturbation n_{t+1} to the log SDF either (i) follows a martingale difference sequence or (ii) features transitory, unconditional-mean-zero disturbances, as is the case in many models.[47] The ν in (A.5) is therefore

$$\nu = -\rho - \widetilde{\gamma}\nu_c. \tag{A.8}$$

[45] It is slightly more general, for example, than the assumption in Backus, Chernov, and Zin (2014, eq. (22)).

[46] In the unit EIS case with $\psi = 1$, $N_{t+1} = \exp(n_{t+1})$ is a martingale, but n_{t+1} is typically not.

[47] This holds in, for example, the Campbell and Cochrane (1999) model, and see Borovička, Hansen, and Scheinkman (2016) for further discussion. If it does not hold, then the ρ in our decomposition should be understood to contain both the time discount rate and any small component arising from $\mathbb{E}[n_{t+1}]$ (e.g., from a Jensen's inequality correction for the continuation utility term in an Epstein–Zin framework). See Appendix A.2 for a discussion of how this affects the r_t^* decomposition in a tractable alternative case.

Similarly, split $\overline{\kappa}(X_t)$ into two parts, $\overline{\kappa}(X_t) = \overline{\kappa}_c(X_t) + \overline{\kappa}_n(X_t)$, where $\overline{\kappa}_c(X_t) = -\widetilde{\gamma}\mu_c(X_t) + \widetilde{\gamma}\nu_c$ and $\overline{\kappa}_n(X_t) = \mu_n(X_t)$. Build up $h_c(X_t)$ and $h_n(X_t)$ accordingly from these $\overline{\kappa}$ functions as in (A.3), and $h(X_t) = h_c(X_t) + h_n(X_t)$. The ε_{t+1} term inherits the remaining martingale-difference components of $-\widetilde{\gamma}(c_{t+1} - c_t)$ and n_{t+1}, as in (A.4).[48] Define these two processes' respective martingale increment terms as $\varepsilon_{c,t+1}$ and $\varepsilon_{n,t+1}$. Finally, define the martingale difference $\widetilde{\varepsilon}_{t+1} = \varepsilon_{t+1} - \varepsilon_{c,t+1} = \varepsilon_{n,t+1}$.

To map the above steps and results to the decomposition in (2), we construct an expanded state vector $\widetilde{X}_t = (c_t, X_t)'$, where X_t is the previous state vector. This expanded state vector still follows a Markov process. With non-zero average consumption growth, $\widetilde{X}_t$ will no longer be stationary, but its differences will be. This expanded state vector will stand in for the state vector used in (2). We can now construct our alternative decomposition. Define

$$f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t) = \widetilde{\gamma}(c_{t+1} - c_t) - (h_n(X_{t+1}) - h_n(X_t)). \tag{A.9}$$

Using (A.5), (A.8), and (A.9), we accordingly have our decomposition

$$m_{t+1} = -\rho - (f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t)) + \widetilde{\varepsilon}_{t+1}, \tag{A.10}$$

where $f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t)$ is a stationary difference and $\widetilde{\varepsilon}_{t+1}$ is a mean-zero martingale difference, as in (2). Note that while $\mathbb{E}_t[f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t)]$ does not depend only on cash-flow growth $c_{t+1} - c_t$ in general, the limiting forward expectation $\lim_{\tau\to\infty} \mathbb{E}_t[f(\widetilde{X}_{t+\tau+1}) - f(\widetilde{X}_{t+\tau})]$ used in the trend real-rate decomposition (5) does:

$$\widetilde{g}_t^* = \lim_{\tau\to\infty} \mathbb{E}_t[f(\widetilde{X}_{t+\tau+1}) - f(\widetilde{X}_{t+\tau})] = \widetilde{\gamma}\mathbb{E}[c_{t+1} - c_t] = \widetilde{\gamma}\nu_c. \tag{A.11}$$

Finally, to complete the characterization of this setting's decomposition, we now allow for unanticipated changes in the economic environment. We accordingly denote the transition distribution describing the Markov process X at date t to be $\mathcal{P}_t$, and assume that this distribution governs X_{t+1} and is expected to govern $X_{t+1+\tau}$ for all $\tau > 0$. There may then be an unanticipated change in the transition distribution to $\mathcal{P}_{t+1}$, at which point this distribution will govern X_{t+2} and will be expected to govern all future $X_{t+2+\tau}$ thereafter. Again defining the expanded state vector $\widetilde{X}_t = (c_t, X_t)'$, the stationary (under $\mathcal{P}_t$) difference $f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t)$ in (A.9), and the now potentially time-varying ρ_t in (A.7), our decomposition (A.10) becomes

$$m_{t+1} = -\rho_t - (f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t)) + \widetilde{\varepsilon}_{t+1}. \tag{A.12}$$

[48] See Borovička, Hansen, and Scheinkman (2016) for explicit characterizations of the martingale components in alternative models.

This now aligns with (2), with $\widetilde{X}_t$ and $\widetilde{\varepsilon}_{t+1}$ here in place of X_t and ε_{t+1} in the text. The limiting expectation $\widetilde{g}_t^*$ as defined in (A.11) may also be time-varying here, as $\widetilde{g}_t^* = \widetilde{\gamma}_t \nu_{c,t}$.

A.1.2 Drifting Steady State

Analogous to Case II in Section 2.3, we now consider an alternative case in which fundamentals (and resulting expected returns and valuation ratios) follow a random walk, or "drifting steady state." Define W_{t+1} as in Appendix A.1.1, and assume that the cash-flow process c and one-dimensional process X are modified from (A.6) to now follow

$$c_{t+1} - c_t = \mu_c + X_t + \sigma_c(X_t) B_c W_{t+1},$$
$$X_{t+1} = X_t + \sigma_x(X_t) B_x W_{t+1}.$$

The volatility term $\sigma_x(X_t)$ may be specified so as to ensure that X remains bounded in L^2 in order to rule out explosive dynamics (or one could assume alternative bounded-martingale dynamics for X), but we do not impose this directly.[49]

In place of (A.7), the log SDF follows

$$m_{t+1} = -\rho_t - \widetilde{\gamma}(c_{t+1} - c_t) + n_{t+1},$$
$$\rho_{t+1} = \rho_t + B_\rho W_{t+1},$$
$$n_{t+1} = \eta_{t+1} - \eta_t,$$

where $\mathbb{E}_t[\eta_{t+1}] = \eta_t$. Defining $\widetilde{X}_t = (c_t, X_t)'$, we again immediately obtain a decomposition of the form (A.12) and (2):

$$m_{t+1} = -\rho_t - (f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t)) + \widetilde{\varepsilon}_{t+1},$$

where $f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t) = \widetilde{\gamma}(c_{t+1} - c_t) = \widetilde{\gamma}(\mu_c + X_t)$ is difference-stationary, and where $\widetilde{\varepsilon}_{t+1} = n_{t+1}$ is a martingale difference.

A.1.3 Stationarity

Finally, analogous to Case III in Section 2.3, we consider a setting building on Appendix A.1.1, without unanticipated breaks but with stationary variation in ρ_t. The stationarity of X_t and ρ_t in this setting will imply that infinite-horizon expectations of the discounting and growth

[49] Instead, the drifting steady state model is intended as a convenient tool to represent persistent fluctuations in fundamentals, rather than being a reasonable candidate model over an infinite horizon.

terms are constant. As a result, we discuss how to redefine long-run expectations in a manner incorporating persistent time variation in these processes.

We start with exactly the same setting as in Appendix A.1.1, with everything through equation (A.6) unchanged. We modify (A.7) slightly to allow for ρ_t to follow stationary Markov dynamics: assume that ρ_t is an element of the state vector X_t in (A.1) and (A.6), and

$$m_{t+1} = -\rho_t - \widetilde{\gamma}(c_{t+1} - c_t) + n_{t+1},$$

with the assumptions on the remaining terms in m_{t+1} unchanged. We also define $\overline{\rho} = \mathbb{E}[\rho_t]$. The above time variation in ρ can be thought of as representing a time-varying subjective discount rate for the representative agent, but it also serves as a stand-in for many other sources of potential variation in the intertemporal marginal rate of substitution. It may arise from demographic changes, a heterogeneous-agents model with time variation in the marginal investor (and differences in investors' personal discount rates), or as discussed later in the appendix, capital flows from foreign investors.

We can then follow nearly exactly the same steps as in Appendix A.1.1 to obtain the following valid decomposition:[50]

$$m_{t+1} = -\rho_t - (f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t)) + \widetilde{\varepsilon}_{t+1},$$

with $f(\widetilde{X}_{t+1}) - f(\widetilde{X}_t)$ defined as in (A.9) and $\widetilde{\varepsilon}_{t+1}$ defined as before (A.9).

In this case, all infinite-horizon expectations defined after the r_t^* decomposition (5) are constant, with $\rho_t^* = \lim_{\tau\to\infty} \mathbb{E}_t[\rho_{t+\tau}] = \overline{\rho}$, $\widetilde{g}_t^* = \lim_{\tau\to\infty} \mathbb{E}_t[f(\widetilde{X}_{t+\tau+1}) - f(\widetilde{X}_{t+\tau})] = \widetilde{\gamma}\nu_c$, and $L_{t,M}^* = \lim_{\tau\to\infty} \mathbb{E}_t[L_{t+\tau}(M_{t+\tau+1})] = \sum_{n=2}^{\infty} \frac{\kappa_n(m_{t+1})}{n!}$, where $\kappa_n(m_{t+1})$ is the n^{th} cumulant of the unconditional log SDF distribution. To formalize long-horizon variation in the terms in our r_t^* decomposition in this context, we redefine the above terms as discounted sums:

$$z_t^* = (1-\delta)\sum_{\tau=0}^{\infty} \delta^\tau \mathbb{E}_t[z_{t+\tau+1}], \tag{A.13}$$

$$z_t \in \{\rho_t, g_t, L_t(M_{t+1}), r_{t+1}^f\},$$

[50] To spell out these steps in further detail: first, with the same definitions and assumptions on ν_c and ν_n as before, the ν in (A.5) becomes $\nu = -\overline{\rho} - \widetilde{\gamma}\nu_c$. We now split $\overline{\kappa}(X_t)$ into three parts, $\overline{\kappa}(X_t) = \overline{\kappa}_\rho(X_t) + \overline{\kappa}_c(X_t) + \overline{\kappa}_n(X_t)$, where $\overline{\kappa}_c(X_t)$ and $\overline{\kappa}_n(X_t)$ are as before, and $\overline{\kappa}_\rho = -\mu_\rho(X_t) + \overline{\rho}$. Build up $h_\rho(X_t)$, $h_c(X_t)$, and $h_n(X_t)$ accordingly from these $\overline{\kappa}$ functions as in (A.3), and $h(X_t) = h_\rho(X_t) + h_c(X_t) + h_n(X_t)$. The ε_{t+1} term inherits the remaining martingale components of $c_{t+1} - c_t$ and n_{t+1} as in (A.4), with no additional term for ρ_t given its stationary Markov dynamics, so we define $\widetilde{\varepsilon}_{t+1}$ as before. The expanded state vector is again $\widetilde{X}_t = (c_t, X_t)'$. The decomposition thus applies as stated.

where $\delta \in (0,1)$ is a loglinearization constant and where $g_{t+\tau+1} \equiv f(\widetilde{X}_{t+\tau+1}) - f(\widetilde{X}_{t+\tau})$. We define the loglinearization constant δ in the appendix for Section 2.3 below, and we discuss how the definition (A.13) maps well to the equity yield decomposition in the stationary case. As $(1-\delta)\sum_{\tau=0}^{\infty}\delta^{\tau} = 1$, equation (A.13) defines the starred long-run terms as weighted averages of all future expected realizations of z_t, as in (13) in the main text.

A.2 Interest-Rate Decomposition

A.2.1 Benchmark Case

The general version of the decomposition for r_t^* provided in equation (5) in Section 2.1 follows immediately from the SDF decomposition derived and discussed in Appendix A.1, along with equations (1) and (3). For the consumption-based version in Section 2.2, the second expression provided in (7) starts from equation (4) and then applies equation (25) of Backus, Chernov, and Martin (2011), which relates the log SDF's cumulants to the consumption-growth cumulants in a power-utility setting according to

$$\kappa_{n,t}(m_{t+1}) = (-\gamma)^n \kappa_{n,t}(g_{t+1}). \tag{A.14}$$

See also equation (8) and the preceding equation in Martin (2013), which provides the same interest-rate expression as in (7). Equation (8) then follows directly from the preceding steps. Again see the previous appendix subsection for details on the definitions of the terms in (5) and (8) given each of our three sets of assumptions on the dynamics of fundamentals.

A.2.2 Extension with Epstein–Zin Preferences

As discussed above equation (A.8) and in detail in footnote 47, if the additional term n_{t+1} in the log SDF specification (A.7) does not feature $\mathbb{E}[n_{t+1}] = 0$, then the ρ (or ρ_t) in our decomposition should be understood to contain both the time discount rate and any small component arising from $\mathbb{E}[n_{t+1}]$. This does not pose serious issues for either the decompositions or for the paper's interpretation of ρ_t: as discussed in Section 3.2, we do not view shifts in ρ_t in the data as likely to be arising purely from changes in aggregate patience among domestic investors. That said, given the use of Epstein–Zin preferences in Section 2.3, we briefly discuss precisely how such preferences affect the interest-rate decomposition in a tractable case.

In place of (6), we follow Epstein and Zin (1989) and set

$$U_t = \left\{(1-\beta_t)C_t^{\frac{1-\gamma}{\theta}} + \beta_t(\mathbb{E}_t[U_{t+1}^{1-\gamma}])^{\frac{1}{\theta}}\right\}^{\frac{\theta}{1-\gamma}},$$

where $\theta \equiv (1-\gamma)/(1-\frac{1}{\psi})$ and where ψ is the elasticity of intertemporal substitution (EIS). We again set $\rho_t = -\log \beta_t$. In a complete-markets setting in which there is a consumption claim whose value is aggregate wealth, the SDF is

$$M_{t+1} = \left(\beta_t \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{1}{\psi}}\right)^{\theta} \left(R^w_{t+1}\right)^{-(1-\theta)}, \tag{A.15}$$

where R^w_{t+1} is the gross return on the wealth portfolio. If we further assume jointly lognormal and homoskedastic C_{t+1} and R^w_{t+1}, the risk-free rate is then

$$r^f_{t+1} = \rho_t + \frac{1}{\psi}\mathbb{E}_t[g_{t+1}] - \frac{\theta}{2\psi^2}\sigma_g^2 - \frac{1-\theta}{2}\sigma_w^2, \tag{A.16}$$

as in Campbell (2018), equation (6.44), where $\sigma_g^2 = \mathrm{Var}_t(g_{t+1})$ and $\sigma_w^2 = \mathrm{Var}_t(r^w_{t+1})$. Given (A.15), the SDF's entropy is

$$L_t(M_{t+1}) = \frac{\theta^2}{2\psi^2}\sigma_g^2 + \frac{(1-\theta)^2}{2}\sigma_w^2 + \frac{\theta(1-\theta)}{\psi}\sigma_{gw}, \tag{A.17}$$

where $\sigma_{gw} = \mathrm{Cov}_t(g_{t+1}, r^w_{t+1})$. Combining this with (A.16), we can write

$$\begin{aligned} r^f_{t+1} &= \rho_t + \frac{1}{\psi}\mathbb{E}_t[g_{t+1}] - L_t(M_{t+1}) + \frac{\theta(\theta-1)}{2}\mathrm{Var}_t\left(\frac{g_{t+1}}{\psi} - r^w_{t+1}\right) \\ &= \tilde{\rho}_t + \frac{1}{\psi}\mathbb{E}_t[g_{t+1}] - L_t(M_{t+1}), \end{aligned} \tag{A.18}$$

where $\tilde{\rho}_t = \rho_t + \frac{\theta(\theta-1)}{2}\mathrm{Var}_t\left(\frac{g_{t+1}}{\psi} - r^w_{t+1}\right)$. As a result, (8) still holds, with $\tilde{\rho}^*_t$ in place of ρ^*_t and with $\frac{1}{\psi}$ in place of γ.

Equivalently, since $L_t(M_{t+1}) = \frac{1}{2}\mathrm{Var}_t\left(\frac{\theta g_{t+1}}{\psi} - (\theta-1)r^w_{t+1}\right)$ from (A.17), we can write

$$r^f_{t+1} = \rho_t + \frac{1}{\psi}\mathbb{E}_t[g_{t+1}] - (1+\omega)L_t(M_{t+1}),$$

$$\text{where} \quad \omega \equiv \frac{\theta(1-\theta)\mathrm{Var}_t\left(\frac{g_{t+1}}{\psi} - r^w_{t+1}\right)}{\mathrm{Var}_t\left(\frac{\theta g_{t+1}}{\psi} - (\theta-1)r^w_{t+1}\right)}$$

is a constant given the homoskedastic setting. Given that our empirical estimation allows for a flexible loading of the risk-free rate on our SDF entropy proxy, this constant of proportionality $1+\omega$ will be incorporated in the empirical versions of the decompositions in such a setting.

As an alternative to the assumption of lognormality after (A.15), assume that g_{t+1} is i.i.d.,

in which case the log SDF becomes $m_{t+1} = -\rho_t - \gamma g_{t+1}$, exactly as in the benchmark case in Section 2.2, so the previous decomposition applies.

A.3 Equity Yields and Duration

Following the main text (and similar to Appendix A.1), we derive our results for equity yields in three different settings. We then briefly discuss how the implications for equity duration follow directly.

A.3.1 Case I (Gordon Growth)

Equity Yields and Risk Premia. Given i.i.d. consumption growth $g_{t+1} = c_{t+1} - c_t$ and $d_{t+1} - d_t = \lambda g_{t+1}$, the setting here mirrors that of Martin (2013, Section 1). He works with a risk premium defined slightly differently than ours: while we use the expected log return and set $rp_t \equiv \mathbb{E}_t[r^{\text{mkt}}_{t+1}] - r^f_{t+1}$, he instead uses the log expected return and considers what we will define as $\widetilde{rp}_t \equiv \log \mathbb{E}_t[R^{\text{mkt}}_{t+1}] - r^f_{t+1}$. By definition of entropy, these two versions of the risk premium differ by

$$\begin{aligned} \widetilde{rp}_t - rp_t &= L_t(R^{\text{mkt}}_{t+1}) \\ &= \sum_{n=2}^{\infty} \frac{\lambda^n \kappa_{n,t}(g_{t+1})}{n!}, \end{aligned} \tag{A.19}$$

where the second line uses that $r^{\text{mkt}}_{t+1} = \lambda g_{t+1} + \text{constant}$ (where the constant is in fact ey) in this i.i.d. setting, and then applies the same relation as in (A.14).

Result 1 of Martin (2013), and in particular equation (7), gives that

$$ey^*_t = r^*_t + \widetilde{rp}^*_t - \sum_{n=1}^{\infty} \frac{\lambda^n \kappa_{n,t}(g_{t+1})}{n!}, \tag{A.20}$$

where we note that the summation in the last term starts with the first cumulant ($n = 1$) rather than the second as in (A.19). This last term is therefore equal to the cumulant-generating function (CGF) for consumption growth, $\mathbf{c}(\vartheta) \equiv \sum_{n=1}^{\infty} \frac{\vartheta^n \kappa_{n,t}(g_{t+1})}{n!}$, evaluated at $\vartheta = \lambda$.

Using (A.19) and (A.20), we have that

$$ey^*_t = r^*_t + rp^*_t - \lambda g^*_t, \tag{A.21}$$

as stated in equation (10) and Result 1.

For the risk-premium expressions in (11), we start from equation (5) of Martin (2013), which gives that $\widetilde{rp}_t^* = \mathbf{c}(\lambda) + \mathbf{c}(-\gamma) - \mathbf{c}(\lambda - \gamma)$, and solve for rp_t^* using (A.19):

$$\begin{aligned} rp_t^* &= \widetilde{rp}_t^* - (\mathbf{c}(\lambda) - \lambda g_t^*) \\ &= \lambda g_t^* + \mathbf{c}(-\gamma) - \mathbf{c}(\lambda - \gamma) \\ &= \sum_{n=2}^{\infty} \frac{(-\gamma)^n \kappa_{n,t}(g_{t+1})}{n!} - \sum_{n=2}^{\infty} \frac{(\lambda - \gamma)^n \kappa_{n,t}(g_{t+1})}{n!} \\ &= L_t(M_{t+1}) - L_t(M_{t+1} R_{t+1}^{\text{mkt}}). \end{aligned} \tag{A.22}$$

The first line uses the definition of the CGF in (A.19); the second substitutes in the $\widetilde{rp}_t^*$ solution above; the third expands the CGFs and uses that the first moments cancel; and the last uses that $m_{t+1} + r_{t+1}^{\text{mkt}} = \text{constant} + (\lambda - \gamma) g_{t+1}$ and applies (A.14). (See also Backus, Chernov, and Martin, 2011, p. 2008, for similar steps.) Both lines of (11) follow directly.

To see that $rp_t = L_t(M_{t+1}) - L_t(M_{t+1} R_{t+1}^{\text{mkt}})$ holds in any no-arbitrage setting, one can follow Backus, Boyarchenko, and Chernov (2018, p. 12): take logs of the pricing equation $\mathbb{E}_t[M_{t+1} R_{t+1}^{\text{mkt}}] = 1$ and use the definition of entropy and equation (1) to obtain

$$\begin{aligned} 0 = \log \mathbb{E}_t[M_{t+1} R_{t+1}^{\text{mkt}}] &= \mathbb{E}_t[m_{t+1} + r_{t+1}^{\text{mkt}}] + L_t(M_{t+1} R_{t+1}^{\text{mkt}}) \\ &= \mathbb{E}_t[r_{t+1}^{\text{mkt}}] - r_{t+1}^f - L_t(M_{t+1}) + L_t(M_{t+1} R_{t+1}^{\text{mkt}}). \end{aligned}$$

Rearranging gives that $rp_t = L_t(M_{t+1}) - L_t(M_{t+1} R_{t+1}^{\text{mkt}})$, as stated.

Using (A.21) and (A.22), along with the interest-rate decomposition derived in the previous appendix sections, Result 1 then follows. We note as well that equations (5) and (7) of Martin (2017) also hold with Epstein–Zin utility, so (A.21) and (A.22) are identical in this case, as stated on page 11.

Risk Shocks. In part (iii) of Result 1, it is stated that equity yields change by $-\frac{\partial rp_t^*}{\partial L_{t,M}^*} + 1$ per unit increase in r_t^* if $\frac{\partial rp_t^*}{\partial L_{t,M}^*}$ is well-defined. The third line of (A.22) shows the need for this qualification: $L_{t,M}^*$ and rp_t^* are both functions of the consumption growth cumulants $\kappa_{n,t}(g_{t+1})$, and there are many potential changes to the different cumulants that generate identical changes in $L_t(M_{t+1})$ but different effects on rp_t^*.[51] In certain settings, though, this partial derivative is well-defined. One such setting is when $\gamma = \lambda$: in this case, $rp_t^* = L_{t,M}^*$,

[51] This motivates our consideration of the average change using the beta of the risk premium with $L_{t,M}^*$, where we use that $\text{Cov}(rp_t^* - L_{t,M}^*, L_{t,M}^*) = \text{Cov}(rp_t^*, L_{t,M}^*) - \text{Var}(L_{t,M}^*)$. When the partial derivative does not exist, one can take $\frac{\partial rp_t^*}{\partial L_{t,M}^*}$ to represent a stand-in for β_L.

so $\frac{\partial rp_t^*}{\partial L_{t,M}^*} = 1$.

Case (i) of the three bond–stock comovement cases described on page 14 is another such setting. This case assumes lognormal growth, so that $\kappa_{n,t}(g_{t+1}) = 0$ for $n > 2$. We therefore have that $rp_t^* = \frac{1}{2}\lambda(2\gamma - \lambda)\mathrm{Var}_t(g_{t+1}) = \frac{\lambda(2\gamma-\lambda)}{\gamma^2}L_{t,M}^*$, so

$$-\frac{\partial rp_t^*}{\partial L_{t,M}^*} + 1 = -\frac{\lambda(2\gamma - \lambda)}{\gamma^2} + 1 = \frac{(\lambda - \gamma)^2}{\gamma^2},$$

which is strictly positive if $\lambda \neq \gamma$. In a lognormal, power-utility setting, the discount-rate effect of an increase in volatility always dominates, so equity prices increase given an increase in volatility as considered here. But the pass-through of interest rates to equity yields (i.e., $-\frac{\partial rp_t^*}{\partial L_{t,M}^*} + 1$ above) is nonetheless strictly below one as long as $(\lambda - \gamma)^2 < \gamma^2$, or equivalently $2\gamma > \lambda$, as stated in the text.

We now consider the other two cases introduced on page 14. Case (ii) is based on the rare-disasters model of Barro (2006). Consumption growth is modeled as $g_{t+1} = g_{1,t+1} + g_{2,t+1}$, where the two components $g_{1,t+1}$ and $g_{2,t+1}$ are independent of each other and independent over time. The first component is normal with $\mathbb{E}_t[g_{1,t+1}] = g_1^*$ and variance σ^2. The second component is a jump component, and we follow Backus, Chernov, and Martin (2011) and assume that it follows a Poisson-normal mixture: a given period has $j \in \mathbb{N}$ jumps with probability $e^{-\omega}\omega^j/j!$ (so ω is the effective jump intensity), and conditional on j, the jump size is normal with mean $-jm$ (with $m > 0$) and variance js^2. This implies (see p. 2002 of Backus, Chernov, and Martin) that the consumption-growth CGF is

$$\mathbf{c}(\vartheta) = \vartheta g_1^* + \frac{1}{2}\vartheta^2\sigma^2 + \omega\left(e^{-\vartheta m + \frac{1}{2}\vartheta^2 s^2} - 1\right).$$

Given that $L_t(\exp(\vartheta g_{t+1})) = \mathbf{c}(\vartheta) - \vartheta\mathbb{E}_t[g_{t+1}]$ and that in this setting $\mathbb{E}_t[g_{t+1}] = g_1^* - \omega m$, we have that

$$\begin{aligned} L_{t,MR}^* &= L_t(\exp((\lambda - \gamma)g_{t+1})) = \mathbf{c}(\lambda - \gamma) - (\lambda - \gamma)(g_1^* - \omega m) \\ &= (\lambda - \gamma)\omega m + \frac{1}{2}(\lambda - \gamma)^2\sigma^2 + \omega\left(e^{-(\lambda-\gamma)m + \frac{1}{2}(\lambda-\gamma)^2 s^2} - 1\right). \end{aligned}$$

Using this,

$$\begin{aligned} ey_t^* &= \rho_t^* + (\gamma - \lambda)g_t^* - L_{t,MR}^* \\ &= \rho_t^* + (\gamma - \lambda)g_1^* - \frac{1}{2}(\lambda - \gamma)^2\sigma^2 - \omega\left(e^{-(\lambda-\gamma)m + \frac{1}{2}(\lambda-\gamma)^2 s^2} - 1\right). \end{aligned}$$

Given a change in the average disaster size m, we therefore have

$$\frac{\partial e y_t^*}{\partial m} = (\lambda - \gamma)\,\omega\, e^{-(\lambda-\gamma)m + \frac{1}{2}(\lambda-\gamma)^2 s^2}$$

This is strictly positive (so valuations go down given an increase in mean disaster size) iff $\lambda > \gamma$.

Meanwhile, following similar steps for the SDF's entropy and plugging into the risk-free rate decomposition,

$$r_t^* = \rho_t^* + \gamma g_1^* - \frac{1}{2}\gamma^2\sigma^2 - \omega\left(e^{\gamma m + \frac{1}{2}(\lambda-\gamma)^2 s^2} - 1\right).$$

As a result,

$$\frac{\partial r_t^*}{\partial m} = -\gamma\,\omega\, e^{\gamma m + \frac{1}{2}(\lambda-\gamma)^2 s^2} < 0,$$

so r_t^* strictly decreases given an increase in mean disaster size. We conclude that such changes induce negative comovement between equity yields and real rates as long as $\gamma < \lambda$.

More generally, for any change in higher ($n \geqslant 2$) moments $\kappa_{n,t}(g_{t+1})$ for n odd, if $\lambda > \gamma$,

$$\frac{\partial L_{t,M}^*}{\partial \kappa_{n,t}(g_{t+1})} = \frac{(-\gamma)^n}{n!} < 0,$$
$$\frac{\partial L_{t,MR}^*}{\partial \kappa_{n,t}(g_{t+1})} = \frac{(\lambda-\gamma)^n}{n!} > 0.$$

Thus, greater negative skewness (i.e., a decrease in $\kappa_{3,t}$) will increase the SDF's entropy $L_{t,M}^*$ but decrease the entropy of the discounted return $L_{t,MR}^*$, and similarly for other higher odd cumulants. Since the risk-free rate decreases in $L_{t,M}^*$ and the equity yield decreases in $L_{t,MR}^*$, these odd-higher-moment shocks will induce negative comovement in the $\gamma < \lambda$ case.

For case (iii) on page 14, with Epstein–Zin utility (and parameters as defined in Appendix A.2.2), using equations (4)–(5) in Martin (2013) and considering a consumption claim ($\lambda = 1$), we can write

$$r_t^* = \rho_t^* + \sum_{n=1}^{\infty} \frac{\kappa_{n,t}(g_{t+1})}{n!}\left((1-1/\theta)(1-\gamma)^n - (-\gamma)^n\right)$$
$$e y_t^* = \rho_t^* + \sum_{n=1}^{\infty} \frac{\kappa_{n,t}(g_{t+1})}{n!}\left(-(1-\gamma)^n/\theta\right).$$

So considering any change in higher ($n \geqslant 2$) moments $\kappa_{n,t}(g_{t+1})$ for n even, we can use the

assumptions $\psi > 1$, $\gamma > 1$, and plug in for θ (and use $\theta < 0$ with those assumptions) to get

$$\frac{\partial r_t^*}{\partial \kappa_{n,t}(g_{t+1})} = \frac{1}{n!}\left(\frac{1}{\psi} - \gamma\right)(1-\gamma)^{n-1} - (-\gamma)^n < 0,$$
$$\frac{\partial ey_t^*}{\partial \kappa_{n,t}(g_{t+1})} = \frac{1}{n!}\left(-(1-\gamma)^n/\theta\right) > 0,$$

so equity yields and risk-free rates move in opposite directions, as stated.

A.3.2 Case II (Drifting Steady State)

If $\mathbb{E}_t[ey_{t+1}] = ey_t \equiv ey_t^*$ (and all underlying fundamentals similarly are martingales), we can follow Gao and Martin (2021, eq (12)–(14)): since $R_{t+1}^{\text{mkt}} = \frac{D_{t+1}+P_{t+1}}{P_t}$, taking logs and expectations yields

$$ey_t^* = \mathbb{E}_t[r_{t+1}^{\text{mkt}}] - \lambda \mathbb{E}_t[g_{t+1}] - \log\left(1 - e^{-ey_t}\right) + \mathbb{E}_t\left[\log\left(1 - e^{-ey_{t+1}}\right)\right].$$

To a first order for ey_{t+1} around its expectation ey_t, the last two terms cancel, giving equation (12) as stated. Given that r_t^* satisfies the same decomposition (as shown in Appendix A.1.2), Result 1 therefore holds as stated.

A.3.3 Case III (Stationarity)

In this case, we apply Result 2 of Gao and Martin (2021) directly to obtain that to a first order around the unconditional expectation $\overline{ey} \equiv \mathbb{E}[ey_t]$,

$$ey_t^* \equiv ey_t = (1-\delta)\sum_{\tau=1}^{\infty} \delta^\tau \mathbb{E}_t[r_{t+\tau+1}^{\text{mkt}} - \lambda g_{t+\tau+1}],$$

where $\delta \equiv e^{-\overline{ey}}$. Therefore, defining rp_t^* (and its underlying components) and g_t^* as in (13), and using the definition for r_t^* and its underlying components from (A.13), the stated decomposition follows. So given the decomposition for r_t^* as shown in Appendix A.1.3, Result 1 again holds as stated.

A.3.4 Equity Duration

The equity duration implications in Section 2.4 for the most part follow from direct application of the previous expressions. For equation (15), we use that $\mathbb{E}_t[D_{t+n}] = D_t e^{n(\lambda g)}$ and then evaluate the resulting series. The remaining derivatives in (16) and Result 2 then follow directly from (14)–(15).

A.3.5 Empirical Entropy Proxy

As stated in Section 3 (see page 18), if the market is growth-optimal and the distribution of log growth is symmetric, then $L^*_{t,M,j} = L^*_{t,R,j}$. To see this, note from (A.14) and (A.19) that

$$\kappa_{n,t}(m_{t+1}) = (-\gamma)^n \kappa_{n,t}(g_{t+1}),$$
$$\kappa_{n,t}(r^{\text{mkt}}_{t+1}) = \lambda^n \kappa_{n,t}(g_{t+1}).$$

Growth optimality requires $\lambda = \gamma$, and a symmetric log growth distribution implies that $\kappa_{n,t}(g_{t+1}) = 0$ for all odd n (aside from $n = 1$, which does not enter into the entropy sum). And for even n, given $\gamma = \lambda$, we have $(-\gamma)^n = \lambda^n$. As a result, $\kappa_{n,t}(m_{t+1}) = \kappa_{n,t}(r^{\text{mkt}}_{t+1})$ for all n, and so $L^*_{t,M,j} = L^*_{t,R,j}$, as stated. This (along with Martin, 2017, Result 3) motivates our use of the squared VIX to proxy for the SDF entropy term $L^*_{t,M}$ in estimating our r^*_t decomposition.

B. Additional Empirical Details and Results

This appendix provides additional measurement details and empirical results as referenced in the main text.

B.1 Measurement Details

Section 3.1 explains most of the paper's data sources and variable definitions. Here, we provide additional details on two aspects of the data mentioned further in the text.

VIX Measurement. The squared VIX is defined for horizon $T - t$ as

$$\text{VIX}^2_{t,T} = \frac{2R^f_{t,T}}{T-t}\Big(\int_0^{F_{t,T}} \frac{\text{put}_{t,T}(K)}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{\text{call}_{t,T}(K)}{K^2} dK\Big),$$

where $F_{t,T}$ is the forward price and $\text{put}_{t,T}(K)$ and $\text{call}_{t,T}(K)$ are prices of European put and call options with strike K expiring at T. To implement this formula, we use a cleaned version of a global panel of index option prices from OptionMetrics. As in the main text, the sample, data filters, and implementation approach are taken from Gandhi, Gormsen, and Lazarus (2023); see that paper for details.

Our options data are available starting in 1990 in the U.S. sample, but the samples for other countries start between 2002 and 2006. To obtain a full sample corresponding to

the forecast data, we project $\text{VIX}^2_{t,j}$ in the available sample onto realized volatility in the country j index return, and then obtain predicted values $\widehat{\text{VIX}}^2_{t,j}$ using the observed volatility for any dates in the sample for which we cannot calculate VIX directly.

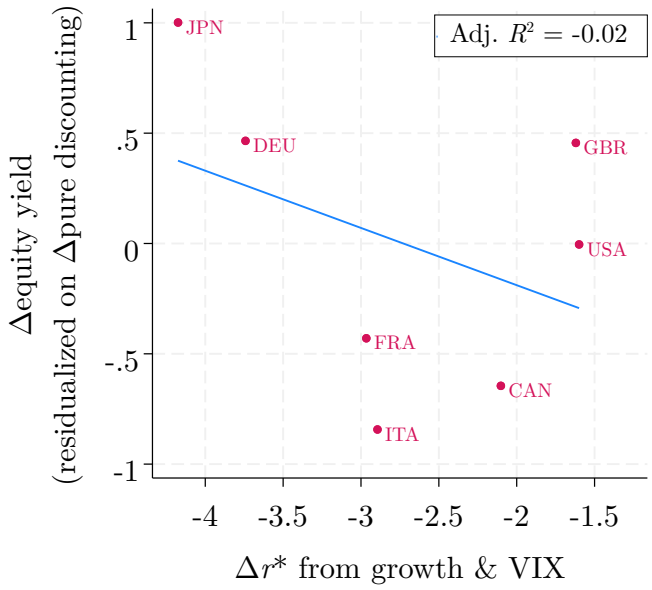
Equity Yields and Data Availability. As in the text, the equity yield $ey_{t,j}$ is measured by starting with the five-year earnings-to-price ratio $\overline{E}_{t-4,t,j}/P_{t,j} = [(E_{t-4,j} + \ldots + E_{t,j})/5]/P_{t,j}$, where earnings and prices are calculated on a value-weighted basis for all available traded stocks in the country. Earnings $E_{t,j}$ are defined as net income for the full calendar year corresponding to date t, while prices $P_{t,j}$ are end-of-period aggregate market capitalizations. We then scale this $\overline{E}_{t-4,t,j}/P_{t,j}$ by 0.5, very close to the unconditional average payout ratio of 0.494 in our post-1990 sample, to obtain our final measure of equity yields $ey_{t,j}$. Algebraically, $ey \equiv \log(1 + D/P) \approx D/P = (D/E) \times (E/P)$. Our unconditional average payout ratio is calculated as the average ratio of five-year-average common dividends to five-year-average net income, to put the dividend and earnings figures in common terms, and the average is across all years and G7 countries starting in 1990.

In some countries, the share of publicly traded companies with available earnings data is low in the early part of our sample. We drop any country–year equity yield observations with such coverage issues. The resulting samples start in 1990 for the U.S. and Canada; 1992 for the U.K.; 1993 for Japan; 1994 for France and Germany; and 1998 for Italy. Any estimated relationship between changes in earnings yields and changes in interest rates uses a country-specific start date consistent with the beginning of this equity yield sample; for example, the r^* difference for Italy in Figure 3 is the difference starting from 1998, consistent with the difference calculated for its equity yields. While all our analyses use data through 2023, the need for full-year net income data to calculate $ey_{t,j}$, along with the fact that our data is only available through part of 2023, means that we measure equity yields through the end of 2022. In all cases where we calculate full-sample differences in the real rate (and its components), we again match this full-sample difference with the actual available sample for our equity data. (Any analysis focusing solely on the real-rate data — for example, the first-stage regression in Table 1 — uses all the data through 2023.)

B.2 Additional Results

See below for additional figures and tables referenced in the main text.

Figure B.1: Residualized Equity Yield Changes vs. Growth and Uncertainty



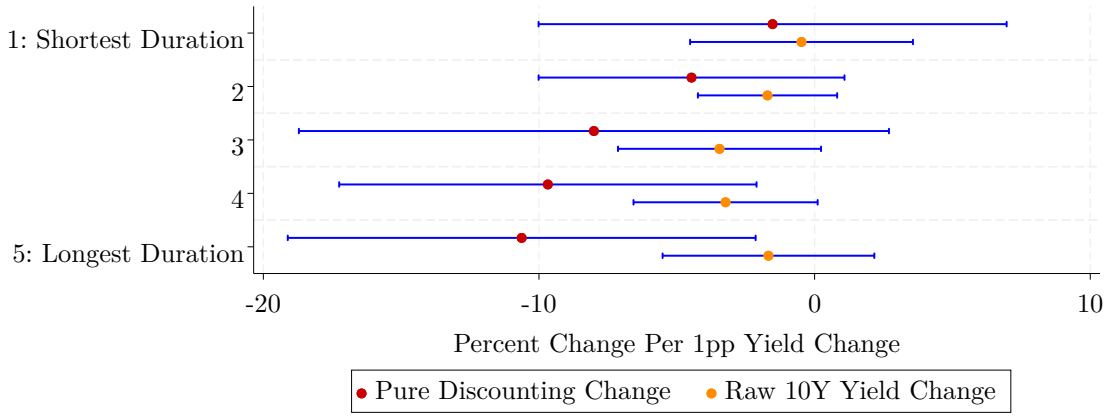
Notes: This figure plots the country-level change in equity yields against changes in interest rates from growth rates and uncertainty, where the equity yield change has now been residualized against the pure discounting change shown in the left panel of Figure 3. The sample is 1990–2023, or the longest available span for the given country. See Figure 3 for additional details.

Table B.1: Regressions for Three-Year Changes in Trend Real Rates

	(1) U.S.	(2) All	(3) All
Change in expected growth $\Delta g^*_{t,j}$	0.5**	0.3**	0.3**
	(0.2)	(0.1)	(0.1)
Change in uncertainty $\Delta \text{VIX}^2_{t,j}$	-4.3**	-0.9	β_j
	(2.1)	(1.8)	
Constant	-0.3***	-0.5***	-0.5***
	(0.1)	(0.1)	(0.1)
Country FEs	✗	✓	✓
Country-Specific $\text{VIX}^2_{t,j}$ Loading	✓	✗	✓
Obs.	74	784	784
R^2	0.17	0.05	0.06
Within R^2	—	0.02	0.04

Notes: This table shows estimated OLS coefficients in the regression (19), along with standard errors in parentheses. In column (1), standard errors are obtained using a block bootstrap. In columns (2)–(3), standard errors are clustered by country and date. Statistical significance at the 10% level, 5% level, and 1% level are denoted by *, **, and ***, respectively. In column (3), the country-specific loadings on the squared VIX, β_j, are statistically significant at the 10% level for 6 of the 12 countries in our sample, and at the 5% level for 3 of the 12 countries (including the U.S.). The sample is 1990–2023.

Figure B.2: Portfolio Exposures to Pure Discount Rate Changes: Global Stocks



Notes: This figure repeats the analysis shown in Figure 5 using the full global sample of stocks. We form duration-sorted portfolios in the international panel following Gormsen and Lazarus (2023), and then we estimate the same regressions as in Figure 5, with country-level fixed effects. The sample is 1990–2023.

Table B.2: Three-Year Return Regressions: Profit-Share Robustness

	(1) U.S.	(2) U.S.	(3) U.S.	(4) U.S.
Δ10y yield	4.19			
	(3.51)			
Δpure discount ($\widehat{\Delta\rho_t^*}$)		-19.1**	-24.6***	-17.9***
		(7.64)	(8.95)	(6.61)
Δexp. growth		-1.49	-15.3	-2.30
		(14.0)	(16.7)	(12.7)
Δexp. profit growth			-0.14	
			(4.02)	
ΔLTG				5.94***
				(2.03)
$\Delta \text{VIX}^2 \times 100$		-3.08**	-4.62***	-2.88***
		(1.33)	(1.46)	(1.09)
Obs.	74	74	58	74
R^2	0.04	0.20	0.47	0.36

Notes: This table replicates the analysis in Table 2 for U.S. data, but with additional predictor variables for other measures related to equity dividend growth. See the notes for Table 2 for details on the estimation and inference, and see Section 4.1 for descriptions of the additional predictor variables.

References

Abel, A. B. (1999): "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics*, 43, 3–33.

Alvarez, F., A. Atkeson, and P. J. Kehoe (2009): "Time-Varying Risk, Interest Rates, and Exchange Rates in General Equilibrium," *Review of Economic Studies*, 76, 851–878.

Alvarez, F. and U. J. Jermann (2005): "Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth," *Econometrica*, 73, 1977–2016.

Andrews, S. and A. Gonçalves (2020): "The Bond, Equity, and Real Estate Term Structures," *Working Paper*.

Asness, C. (2022): "Is Value Just an Interest Rate Bet?" *AQR Research Note*.

Atkeson, A. G., J. Heathcote, and F. Perri (2024): "The End of Privilege: A Reexamination of the Net Foreign Asset Position of the United States," *Forthcoming, American Economic Review*.

Auclert, A., H. Malmberg, F. Martenet, and M. Rognlie (2024): "Demographics, Wealth, and Global Imbalances in the Twenty-First Century," *NBER Working Paper 29161*.

Backus, D., N. Boyarchenko, and M. Chernov (2018): "Term Structures of Asset Prices and Returns," *Journal of Financial Economics*, 129, 1–23.

Backus, D., M. Chernov, and I. Martin (2011): "Disasters Implied by Equity Index Options," *Journal of Finance*, 66, 1969–2012.

Backus, D., M. Chernov, and S. Zin (2014): "Sources of Entropy in Representative Agent Models," *Journal of Finance*, 69, 51–99.

Backus, D. K., S. Foresi, and C. I. Telmer (2001): "Affine Term Structure Models and the Forward Premium Anomaly," *Journal of Finance*, 56, 279–304.

Baele, L., G. Bekaert, and K. Inghelbrecht (2010): "The Determinants of Stock and Bond Return Comovements," *Review of Financial Studies*, 23, 2374–2428.

Bansal, R. and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

Barro, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121, 823–866.

Barsky, R. B. (1989): "Why Don't the Prices of Stocks and Bonds Move Together?" *American Economic Review*, 79, 1132–1145.

Bauer, M. D. and G. D. Rudebusch (2020): "Interest Rates under Falling Stars," *American Economic Review*, 110, 1316–1354.

Bauer, M. D. and E. T. Swanson (2023a): "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual*, 37, 87–155.

——— (2023b): "An Alternative Explanation for the 'Fed Information Effect'," *American Economic Review*, 113, 664–700.

van Binsbergen, J. H. (2024): "Duration-Based Stock Valuation: Reassessing Stock Market Performance and Volatility," *Forthcoming, Journal of Financial Economics*.

van Binsbergen, J. H., L. Ma, and M. Schwert (2024): "The Factor Multiverse: The Role of Interest Rates in Factor Discovery," *Working Paper*.

Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer (2024): "Belief Overreaction and Stock Market Puzzles," *Journal of Political Economy*, 132, 1450–1484.

Borovička, J., L. P. Hansen, and J. A. Scheinkman (2016): "Misspecified Recovery," *Journal of Finance*, 71, 2493–2544.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2008): "An Equilibrium Model of 'Global Imbalances' and Low Interest Rates," *American Economic Review*, 98, 358–393.

Camanho, N., H. Hau, and H. Rey (2022): "Global Portfolio Rebalancing and Exchange Rates," *Review of Financial Studies*, 35, 5228–5274.

Campbell, J. Y. (1986): "Bond and Stock Returns in a Simple Exchange Model," *Quarterly Journal of Economics*, 101, 785–803.

——— (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

——— (2008): "Viewpoint: Estimating the equity premium," *Canadian Journal of Economics*, 41, 1–21.

——— (2018): *Financial Decisions and Markets: A Course in Asset Pricing*, Princeton: Princeton University Press.

Campbell, J. Y. and J. Ammer (1993): "What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns," *Journal of Finance*, 48, 3–37.

Campbell, J. Y. and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

Campbell, J. Y., C. Pflueger, and L. M. Viceira (2020): "Macroeconomic Drivers of Bond and Equity Risks," *Journal of Political Economy*, 128, 3148–3185.

Campbell, J. Y. and S. B. Thompson (2008): "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *Review of Financial Studies*, 21, 1509–1531.

Catherine, S., M. Miller, J. D. Paron, and N. Sarin (2023): "Interest-Rate Risk and Household Portfolios," *Working Paper*.

Chernov, M., L. A. Lochstoer, and D. Song (2023): "The Real Channel for Nominal Bond-Stock Puzzles," *Working Paper*.

Cochrane, J. H. (2018): "A Brief Parable of Over-Differencing," *Unpublished Note*.

Constantinides, G. M. and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104, 219–240.

David, A. and P. Veronesi (2013): "What Ties Return Volatilities to Price Valuations and Fundamentals?" *Journal of Political Economy*, 121, 682–746.

——— (2016): "The Economics of the Comovement of Stocks and Bonds," in *Handbook of Fixed-Income Securities*, ed. by P. Veronesi, Hoboken, NJ: Wiley, chap. 15, 313–326, 1st ed.

Dechow, P. M., R. G. Sloan, and M. T. Soliman (2004): "Implied Equity Duration: A New Measure of Equity Security Risk," *Review of Accounting Studies*, 9, 197–228.

Di Tella, S., B. Hébert, P. Kurlat, and Q. Wang (2024): "The Zero-Beta Interest Rate," *Working Paper*.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

Fagereng, A., M. Gomez, É. Gouin-Bonenfant, M. Holm, B. Moll, and G. Natvik (2024): "Asset-Price Redistribution," *Forthcoming, Journal of Political Economy*.

Fama, E. F. and K. R. French (1993): "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

Farhi, E. and F. Gourio (2018): "Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia," *Brookings Papers on Economic Activity*, 147–250.

Gandhi, M., N. J. Gormsen, and E. Lazarus (2023): "Forward Return Expectations," *NBER Working Paper 31687*.

Gao, C. and I. W. R. Martin (2021): "Volatility, Valuation Ratios, and Bubbles: An Empirical Measure of Market Sentiment," *Journal of Finance*, 76, 3211–3254.

Gormsen, N. J. and E. Lazarus (2023): "Duration-Driven Returns," *Journal of Finance*, 78, 1393–1447.

Greenwald, D. L., M. Leombroni, H. Lustig, and S. Van Nieuwerburgh (2023): "Financial and Total Wealth Inequality with Declining Interest Rates," *Working Paper*.

Greenwald, D. L., M. Lettau, and S. C. Ludvigson (2025): "How the Wealth Was Won: Factor Shares as Market Fundamentals," *Journal of Political Economy*, 133, 1083–1132.

Griliches, Z. and J. A. Hausman (1986): "Errors in Variables in Panel Data," *Journal of Econometrics*, 31, 93–118.

Hansen, L. P. (2012): "Dynamic Valuation Decomposition Within Stochastic Economies," *Econometrica*, 80, 911–967.

——— (2019): "Additive Functionals," *Lecture Notes*.

Hansen, L. P. and T. J. Sargent (2022): "Risk, Uncertainty, and Value," *Manuscript, University of Chicago and New York University*.

Hanson, S. G. and J. C. Stein (2015): "Monetary Policy and Long-Term Real Rates," *Journal of Financial Economics*, 115, 429–448.

Hassan, T., T. M. Mertens, and J. Wang (2024): "A Currency Premium Puzzle," *Federal Reserve Bank of San Francisco Working Paper 2024-32*.

Itskhoki, O. and D. Mukhin (2021): "Exchange Rate Disconnect in General Equilibrium," *Journal of Political Economy*, 129, 2183–2232.

Jiang, Z., A. Krishnamurthy, and H. Lustig (2024): "Implications of Asset Market Data for Equilibrium Models of Exchange Rates," *Working Paper*.

Kekre, R. and M. Lenel (2024): "Exchange Rates, Natural Rates, and the Price of Risk," *Working Paper*.

Knox, B. and A. Vissing-Jorgensen (2024): "A Stock Return Decomposition Using Observables," *Working Paper*.

Kremens, L., I. W. R. Martin, and L. Varela (2024): "Long-Horizon Exchange Rate Expectations," *Working Paper*.

Kroen, T., E. Liu, A. R. Mian, and A. Sufi (2024): "Falling Rates and Rising Superstars," *Working Paper*.

Laarits, T. (2022): "Precautionary Savings and the Stock-Bond Covariance," *Working Paper*.

Lazarus, E., D. J. Lewis, J. H. Stock, and M. W. Watson (2018): "HAR Inference: Recommendations for Practice," *Journal of Business & Economic Statistics*, 36, 541–559.

Lettau, M. and J. A. Wachter (2007): "Why is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium," *Journal of Finance*, 62, 55–92.

Lustig, H. and A. Verdelhan (2019): "Does Incomplete Spanning in International Financial Markets Help to Explain Exchange Rates?" *American Economic Review*, 109, 2208–2244.

Maloney, T. and T. J. Moskowitz (2021): "Value and Interest Rates: Are Rates to Blame for Value's Torments?" *Journal of Portfolio Management*, 47, 65–87.

Martin, I. (2017): "What Is the Expected Return on the Market?" *Quarterly Journal of Economics*, 132, 367–433.

Martin, I. W. R. (2013): "Consumption-Based Asset Pricing with Higher Cumulants," *Review of Economic Studies*, 80, 745–773.

Nagel, S. and Z. Xu (2022): "Asset Pricing with Fading Memory," *Review of Financial Studies*, 35, 2190–2245.

——— (2024): "Movements in Yields, not the Equity Premium: Bernanke-Kuttner Redux," *NBER Working Paper 32884*.

Nakamura, E. and J. Steinsson (2018): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *Quarterly Journal of Economics*, 133, 1283–1330.

Pavlova, A. and R. Rigobon (2007): "Asset Prices and Exchange Rates," *Review of Financial Studies*, 20, 1139–1180.

Song, D. (2017): "Bond Market Exposures to Macroeconomic and Monetary Policy Risks," *Review of Financial Studies*, 30, 2761–2817.