Equity Duration and Interest Rates*

Niels Joachim Gormsen and Eben Lazarus

March 2025

Abstract

We study how changes in interest rates affect discount rates and prices in equity markets. While stocks have long-duration cash flows, this need not imply that they are highly sensitive to arbitrary interest-rate changes: rates can fall for multiple structural reasons, each of which affects equity differently. We provide a simple theoretical decomposition under which any change in trend real rates can be split into three components: a change in expected growth, a change in uncertainty, and a pure discounting shock. Only the pure discounting term transmits one for one from interest rates to equity yields. Implementing our decomposition with a global panel of growth expectations and asset prices, we find: (i) a weak unconditional relation between stock valuations and trend real rates, but (ii) a very strong relation between valuations and the pure discounting component of rates, with pure discount rate changes explaining over 80% of the cross-country variation in stock valuations since 1990. In the U.S. data, we find that only 35% of the decline in interest rates is attributable to the pure discount 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, E43, E44

*Preliminary draft. We are grateful to Michael Bauer, Martin Beraja, Darrell Duffie, Matthieu Gomez, Martin Lettau, Dmitry Livdan, Hanno Lustig, Peter Maxted, Thomas Mertens, Emi Nakamura, David Sraer, Jón Steinsson, Johan Walden, and seminar participants at the SF Fed Macro–Finance Workshop and UC Berkeley for very helpful comments. Gormsen thanks the Fama-Miller Center, the Fama Faculty Fellowship, and the Asness Junior Faculty Fellowship for financial support. Gormsen is at the University of Chicago Booth School of Business and NBER, and Lazarus is at the Haas School of Business, University of California, Berkeley. Contact: niels.gormsen@chicagobooth.edu and lazarus@berkeley.edu.

1. Introduction

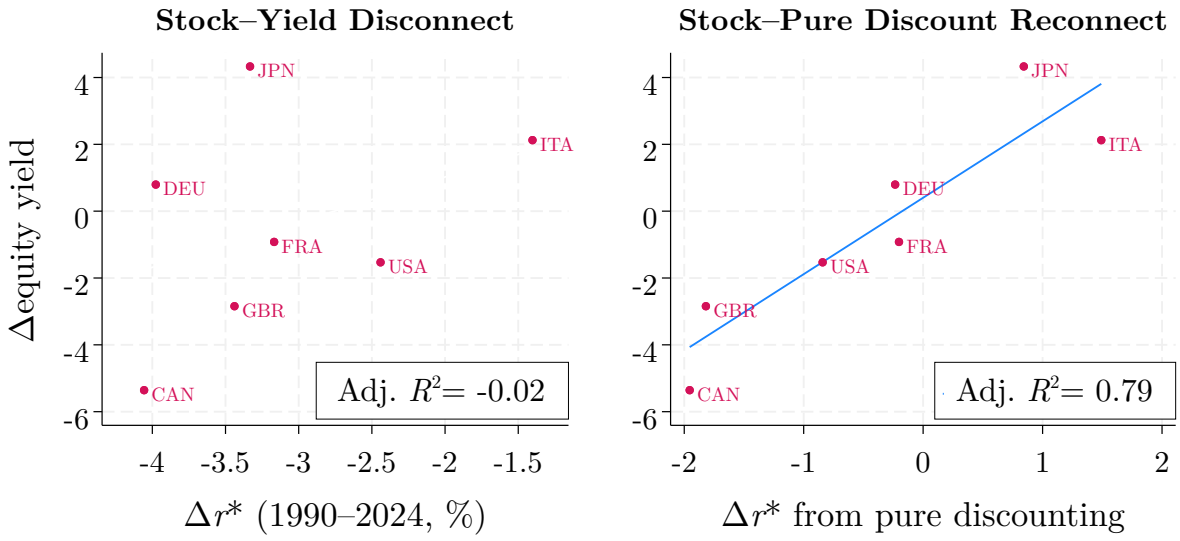
Advanced economies' long-term interest rates have declined significantly in recent decades. How do such changes in rates transmit to equity valuations? A potentially tempting line of reasoning is to assume that equity discount rates move one for one with interest rates; given that equity is a long-duration 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 — measured as the earnings-to-price ratio — 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-rate changes do not arise by happenstance; instead, 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, decrease because of a decrease in expected growth rates in the economy, which — keeping all else constant — will decrease stock prices and mute the effects of the decline in discount 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 trend real rates is

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 1990–2024, 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.

therefore key 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 account for the underlying structural drivers. The punchline of our empirical implementation is that most of the secular change 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 explain a substantial part of the fluctuations in stock prices at shorter 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 stock returns. 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

[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 the realized data.

decreases are important for explaining the strong performance of U.S. stocks and bonds in 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.

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:

1. 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. Our theory suggests that 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.

2. 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 challenge this notion empirically, at least in part: our

[2]Similarly, Andrews and Gonçalves (2020) estimate a decreasing term structure of dividend risk premia relative to maturity-matched zero-coupon 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.

findings imply 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 inequality was non-mechanical.

3. 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 results clarify 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.

4. 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

[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 value-tilted portfolio returns lower."

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, 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

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

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 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. In Section 4, we analyze additional implications of our findings for asset prices in recent decades and related literature. Section 5 discusses our results and concludes.

[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 (2024) argue that increased profit shares are a key driver of equity-price increases in the U.S. data (and attribute relatively little of the increase to interest-rate changes). While their model-based estimation differs from ours, their results and ours are in principle consistent. An increase in profit shares will increase both equity prices and earnings; their main focus is on the price increase, while ours is on valuation ratios.

[7]Others, including Pavlova and Rigobon (2007) and Camanho, Hau, and Rey (2022), additionally consider equity markets and capital flows. 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. 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.

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, following Theorem 3.1 of Hansen (2012),[9] that the log SDF can be additively decomposed as

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

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

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

[9] To be slightly more precise, denoting $M_{t+1} = \widetilde{M}_{t+1}/\widetilde{M}_t$, we assume that $\widetilde{m}_{t+1} = \log \widetilde{M}_{t+1}$ satisfies an additive decomposition in the sense of Hansen's Theorem 3.1. See Hansen (2012) for formal details and a discussion of primitive sufficient conditions under which the assumption holds.

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 labels on the first two terms correspond to the interpretations introduced above. For the labeling of $L_t(M_{t+1})$ as corresponding to an uncertainty or precautionary savings term, note that by the definition of entropy, $L_t(M_{t+1}) = \sum_{n=2}^{\infty} \kappa_{n,t}(m_{t+1})/n!$, 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.[10] Equation (3) directly implies that the trend real rate satisfies

$$r^*_t = \rho^*_t + \widetilde{g}^*_t - L^*_{t,M}, \tag{4}$$

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 (4), we now consider a more standard consumption-based framework. We assume an endowment economy in which a representative

[10] 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, Chernov, and Zin (2014).

agent has power utility over consumption,

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

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.[11] Denoting log consumption growth by $g_{t+1} = c_{t+1} - c_t$, the log SDF is then $m_{t+1} = -\rho_t - \gamma g_{t+1}$. Plugging this into (1),

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

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 Martin (2013) or Backus, Chernov, and Martin (2011). The trend real rate is therefore

$$r_t^* = \rho_t^* + \gamma g_t^* - L_{t,M}^*, \tag{5}$$

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 $L_{t,M}^* = \lim_{\tau\to\infty} \mathbb{E}_t[L_{t+\tau}(M_{t+\tau+1})]$. Thus exactly as in (4), 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 and the evolution of their moments over time. Our framework here will be quite standard. We will 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

[11] We could generalize to Epstein–Zin preferences or time-varying risk aversion, but the basic takeaways — that ρ_t will affect real rates and equity yields one for one, while growth rates and entropy will have ambiguous or offsetting effects on rates and equity — will be unchanged, so we maintain our simple approach and see how far it takes us empirically.

settings.

For equity cash flows, we follow Campbell (1986) and Abel (1999) and model equity as a levered claim to consumption, paying dividends $D_t = C_t^{\lambda}$. In addition to serving as a reduced-form leverage parameter, λ can also be thought of as a stand-in for the capital (or profit) share of income: an increase in λ without a corresponding change in consumption, for example, represents a case in which all output growth accrues to capital owners. We assume that λ is constant for now; as we will see, such an assumption works well to explain the long-term trends observed in the data.[12]

For the remaining stochastic processes, we consider three cases, each of which delivers approximately identical intuition for equity valuations.

Case I (Martin, 2013): Assume that log output growth $g_{t+1} = c_{t+1} - c_t$ is independently and identically distributed over time, with arbitrary distribution. We view this as a reasonable approximation given our focus on relatively long horizons (over which outcomes are likely to be roughly conditionally i.i.d.), though this conditional i.i.d. assumption will mean that any changes in consumption moments or preference parameters are assumed to be unexpected permanent shocks. In particular, define $g_t^* = \mathbb{E}_t[g_{t+1}] = \mathbb{E}_t[g_{t+\tau}]$ for all $\tau \geqslant 1$. Further, assume the higher-order growth-rate cumulants $\kappa_{n,t+\tau}(g_{t+\tau+1})$ are taken to be constant for all τ. The time indexes are included to allow for unanticipated shocks to these values, which are then taken to be permanent thereafter. We similarly allow for unanticipated shocks to the preference parameter $\rho_t^* = \rho_t$.

The equity yield is defined as

$$ey_t \equiv \log(1 + D_t/P_t), \tag{6}$$

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 (6) because it puts the equity yield in equivalent terms as the log real rate. In addition, this definition allows us to use Result 1 from Martin (2013) directly to obtain the following Gordon growth formula for equity valuations (in which we use the i.i.d. assumption to obtain that all relevant variables are equal to their conditional steady-state values):

$$ey_t^* = r_t^* + rp_t^* - \lambda g_t^*, \tag{7}$$

[12] In future iterations of this project, we plan to relax this assumption. Our survey data includes profit growth forecasts in addition to output growth forecasts, allowing for an empirical proxy for long-term expected profit share changes. In the U.S. data, forecasted profit growth has in fact decreased by more than forecasted output growth, indicating that our empirical results are likely to be robust for this case, but we plan to conduct this analysis more systematically in the full panel of countries.

where $rp_t^* = rp_t \equiv \mu_t - r_t^f = \log \mathbb{E}_t[R_{t+1}^{\text{mkt}}] - r_t^f$ is the log equity premium. This premium solves to

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

The first line again applies Result 1 from Martin (2013), and the second then follows using the relevant definitions.[13] If $\lambda = \gamma$, then the market is the growth-optimal portfolio, $M_{t+1}R_{t+1}^{\text{mkt}} = 1$, and $rp_t^* = L_{t,M}^* + L_{t,R}^*$.[14] If, in addition, the log growth distribution is symmetric so that $\kappa_{n,t}(g_{t+1}) = 0$ for odd n, then $L_{t,M}^* = L_{t,R}^*$ and $rp_t^* = 2L_{t,M}^*$.

Alternatively, for arbitrary λ and γ, if growth is lognormal so that $\kappa_{n,t}(g_{t+1}) = 0$ for all $n > 2$, then $rp_t^* = \lambda\gamma \text{Var}_t(g_{t+1}) = \frac{2\lambda}{\gamma} L_{t,M}^*$. This is a useful benchmark case for obtaining particularly stark results. Using this along with along with (5) in (7), 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*

$$\begin{aligned} r_t^* &= \rho_t^* + \gamma g_t^* - L_{t,M}^*, \\ ey_t^* &= \rho_t^* + (\gamma - \lambda) g_t^* + L_{t,R}^* - L_{t,MR}^* \\ &= \rho_t^* + (\gamma - \lambda) g_t^* + \frac{2\lambda - \gamma}{\gamma} L_{t,M}^*. \qquad \textit{(lognormal)} \end{aligned}$$

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 will have different implications for equity valuations:

(i) **Pure discounting shocks:** Bonds and equity co-move perfectly, with ey_t^* changing by 1% for each 1% change in r_t^*.

(ii) **Growth rate shocks:** Equity yields change by $\frac{\gamma - \lambda}{\gamma}$ for each 1% change in r_t^*. In the case with $\gamma = \lambda$ (e.g., with log utility and an unlevered consumption claim), ey_t^* is unaffected by changes in r_t^* induced by growth-rate shocks.

(iii) **Risk shocks:** Equity yields change by $\frac{\gamma - 2\lambda}{\gamma}$ for each 1% change in r_t^* (in the lognormal case). If $2\lambda > \gamma$, this shock to $L_{t,M}^*$ induces negative comovement between bonds

[13] The second line mirrors Backus, Boyarchenko, and Chernov (2018, p. 12–13).

[14] 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: $L_t(M_{t+1}) = \mathbb{E}_t[r_{t+1}^{\text{mkt}}] - r_t^f$, and the right side differs from rp_t^* by $L_t(R_{t+1}^{\text{mkt}})$.

and stocks. More generally, risk shocks generate negative comovement as long as $\text{Cov}(L^*_{t,R} - L^*_{t,MR}, L^*_{t,M}) \geqslant 0$.

As a result, 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 after presenting our empirical results.

For changes in growth rates, 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 given that it decreases both discount rates and growth rates.[15] The typical case is generally $\gamma \leqslant \lambda$, so that a decrease in growth rates also decreases equity valuations (corresponding to a higher equity yield). This would mean growth-rate changes induce weakly negative comovement between bonds and stocks. For changes in risk, there are offsetting effects on the risk-free rate r^*_t and the risk premium rp^*_t. These changes may approximately offset or may again cause equities to move in the opposite direction of bonds. We examine these relations further in our empirical estimation.

Result 1 is our main decomposition for trend real rates and equity valuations. We now briefly discuss how similar decompositions carry 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) and as in Gao and Martin (2021). 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 $ey_t = \mathbb{E}_t[r^{\text{mkt}}_{t+1} - \lambda g_{t+1}]$.[16] Defining $rp_t \equiv \mathbb{E}_t[r^{\text{mkt}}_{t+1}] - r^f_t$ in this case, this implies that

$$ey^*_t = ey_t = \mathbb{E}_t[r^{\text{mkt}}_{t+1} - \lambda g_{t+1}] = \mathbb{E}_t[r^{\text{mkt}}_{t+2} - \lambda g_{t+2}] = \ldots = r^*_t + rp^*_t - \lambda g^*_t,$$

[15] The equity yield decomposition in Result 1 does not separate between discount rates and growth rates in the same way that a Campbell–Shiller decomposition does. Instead, it collects terms in such a way that the g^*_t term, for example, contains both the direct cash-flow effect (λg^*_t) and the discount-rate effect (γg^*_t). We do so because of our desire to decompose risk-free discount rates into underlying structural components, rather than composite terms.

[16] This follows the derivation used for Gao and Martin (2021, eq. (14)).

assuming the individual limiting values exist as $t + \tau \to \infty$. In addition, r_t^* satisfies exactly the same decomposition as in Result 1. As a result, we are in nearly an identical situation as in Result 1, with the following decomposition for the equity dividend yield:

$$ey_t^* = \rho_t^* + (\gamma - \lambda)g_t^* + (rp_t^* - L_{t,M}^*).$$

The only slight distinction is in the last term: for simplicity, we keep rp_t^* rather than substituting it out for entropy terms. But the main takeaways are exactly the same. Only pure discounting shocks transmit one-for-one from rates to equities. Growth rate shocks roughly offset for $\gamma \approx \lambda$, and risk shocks have ambiguous effects as well (with likely offsetting effects on the risk-free rate and risk premium, potentially inducing negative comovement between rates and equities).

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

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

This is 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, and that there are no unanticipated permanent shocks. This case does not admit permanent changes to real rates or valuations. But reinterpreting r_t^*, rp_t^* and g_t^* as Campbell-Shiller–type discounted sums — for example, $r_t^* = (1-\delta)\sum_{\tau=1}^{\infty}\delta^\tau \mathbb{E}_t[r_{t+\tau}^f]$ — generates an approximate Gordon growth formula $ey_t = r_t^* + rp_t^* - \lambda g_t^*$, again following Gao and Martin (2021). We thus obtain very closely analogous results relating equity yields to changes in each of the long-term components underlying interest rates.

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:

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

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

$$\mathcal{D} \equiv \sum_{n=1}^{\infty} n \frac{e^{-n(\mu^*)}\mathbb{E}_t[D_{t+n}]}{P} = \frac{1}{1 - e^{-(\mu^* - \lambda g^*)}} \approx \frac{1}{\mu^* - \lambda g^*}. \tag{8}$$

With this definition and solution, equity duration is equivalent to the following two sensitivities.

(i) **Price sensitivity to the equity discount rate μ^*:**

$$\mathcal{D} = -\frac{\partial \log P}{\partial \mu^*} = \frac{1}{1 - e^{-(\mu^* - \lambda g^*)}} \approx \frac{1}{\mu^* - \lambda g^*}. \tag{9}$$

This is by construction. As in the standard yield-curve context, duration describes percent price sensitivity to a parallel shift in log discount rates.

(ii) **Price sensitivity to the pure discount rate ρ^*:** $\frac{\partial \mu^*}{\partial \rho^*} = 1$, so

$$-\frac{\partial \log P}{\partial \rho^*} = -\frac{\partial \log P}{\partial \mu^*} = \mathcal{D}. \tag{10}$$

By contrast, equity duration is not price sensitivity to a change in r^* arising from the other two components of the interest-rate decomposition. For example, in the case with log utility ($\gamma = 1$) in which equity is an unlevered consumption claim ($\lambda = 1$), the price sensitivity of equity to a change in interest rates due to g^* is exactly zero. Confusion arising in recent literature (as discussed, for example, in footnote 4) suggests that making this distinction clearly is useful.

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 (8), requires a very long sample for statistical precision and is inherently backward-looking. Measurement using price sensitivity to μ^*, as in (9), 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 (4)): $r_t^* = \rho_t^* + \gamma g_t^* - L_{t,M}^*$, where ρ_t^* is the pure discounting term (or rate of time preference), g_t^* is long-term expected output growth, $L_{t,M}^*$ 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.[17] 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 2024, 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

[17] 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/.

inflation forecast at the same horizon.[18] 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 relationship) is of course possible, two points are worth noting. First, a sizable share of the forecasters work at large financial institutions that play a key role in trading and pricing assets. Second, our main exercise of interest will be to use our decomposition to measure transmission to traded equity prices. These equity prices will therefore allow for an out-of-sample validation of our measurement for interest rates, by testing whether our measured real-rate 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 derived after equation (2.3), 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.[19] These imply that setting $L^*_{t,M,j} \propto \text{VIX}^2_{t,j}$ is likely to provide a reasonable approximation, and the constant of proportionality will be implicitly estimated in our regressions for real rates below.

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 (see footnote 19) 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. 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

[18] This approach roughly matches the statistical estimates 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.

[19] The VIX is defined for horizon $T-t$ as $\text{VIX}^2_{t,T} = \frac{2R^f_{t,T}}{T-t}\left(\int_0^{F_{t,T}} \frac{\text{put}_{t,T}(K)}{K^2}dK + \int_{F_{t,T}}^{\infty} \frac{\text{call}_{t,T}(K)}{K^2}dK\right)$, 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.

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.

For equity prices and valuation ratios, we use a value-weighted index for each country based on data from CRSP and Compustat (via the XpressFeed global database). We calculate equity yields $ey_{t,j}$ as 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.[20] Note that our use of an earnings yield for $ey_{t,j}$ differs slightly from the dividend yield considered in our theory. We do so to sidestep the issue that dividend payout ratios have declined in recent decades, which may induce decreases in $dp_{t,j}$ without any changes in forward-looking growth rates. That said, the cross-country correlation of differences in $dp_{t,j}$ and $ey_{t,j}$ from the start to the end of the sample is above 0.8, indicating that our results are not overly sensitive to the use of either definition of equity yields. The only distinction is that the earnings yield is equal to a scaled-up version of the equity yield considered in our theory, with a scaling factor of E/D. In some countries,

Finally, when conducting our cross-sectional analyses, we use returns on duration-sorted portfolios via Gormsen and Lazarus (2023). That paper starts from analyst forecasts of long-term expected earnings growth (via IBES), or LTG. Given that analysts only cover a subset of firms, we then project LTG onto a set of five firm characteristics, and construct a predicted $\widehat{\text{LTG}}$ for all firms. We then construct portfolios sorted by this predicted LTG and calculate their returns. See Gormsen and Lazarus (2023) for details. We also obtain data on size-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

$$r^*_{t,j} = \rho_0 + \gamma\, g^*_{t,j} + \beta\, \text{VIX}^2_{t,j} + \Gamma_j + \varepsilon_{t,j}, \tag{11}$$

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

[20] 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 from 1998 to 2024, consistent with the difference calculated for its equity yields.

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 (11), 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 significantly significant at the 1% level for 9 of the 12 countries in our sample. The sample is 1990–2024, or the longest available span for the given country.

cases in which sovereign credit risk affects a country's r^*_t.[21] 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

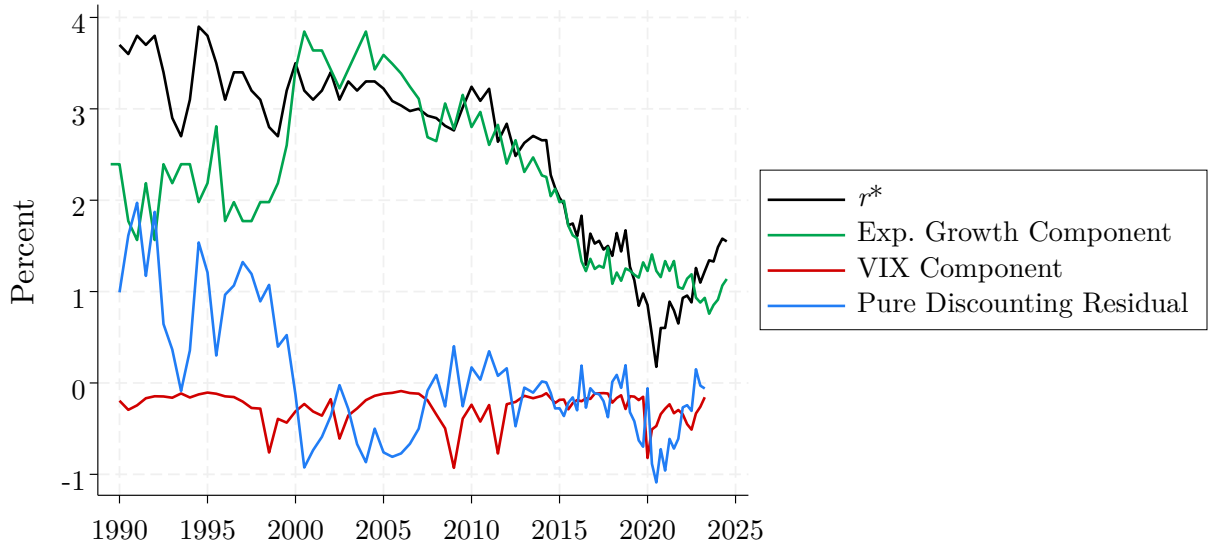
$$\widehat{\rho}^*_{t,j} = \widehat{\rho}_0 + \widehat{\Gamma}_j + \widehat{\varepsilon}_{t,j}. \tag{12}$$

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 (11) 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

[21] 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



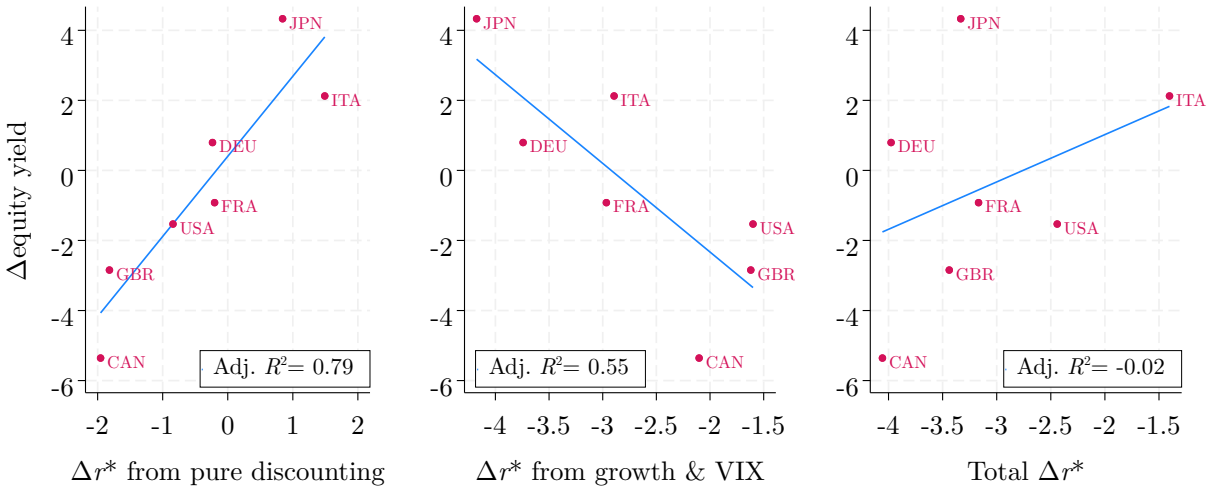
Notes: This figure shows the U.S. trend real rate $r^*_{t,j}$ and its components over time, estimated using (11)–(12) following the main specification in column (3) of Table 1. For readability, the expected growth component is shifted down by 3 percentage points ($\widehat{\gamma} g^*_{t,j} - 3$), and 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 present 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 residual 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 (11)–(12) 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 pure discounting term; the rightmost figure plots changes in equity yields against changes real rates themselves. The sample is 1990–2024, or the longest available span for the given country. For countries for which we can only measure equity yields starting after 1990 (see footnote 20), 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 (11)–(12), 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.

In addition, 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 two percentage points for every one-percentage-point decrease in the pure discounting term. Since we are looking at earnings-to-price ratios — and not dividend-to-price ratios — the prediction is not a one-to-one relation between equity yields and the pure discounting term. Instead, the change in equity earnings yields E/P should be equal to one divided by the payout ratio, which is indeed very close to two given payout ratios of 0.5.

Intuitively, a decrease in earnings of two dollars decreases dividends only by one dollar, so it is associated with a less than one-to-one pricing impact.[22] We could have equivalently plotted dividend price ratios on the left hand side, for which the slope is indeed very close to one.

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.

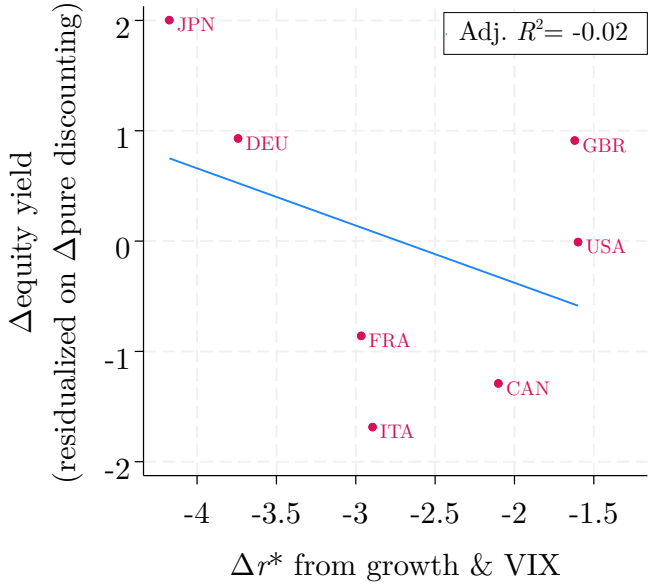
The middle plot in Figure 3 illustrates the relation between earnings yields and the change in interest rates induced by changes in expected growth rates and uncertainty, taken together. As expected, we find that valuation ratios have dropped in countries where interest rates have dropped because of 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 interpret this set of findings. 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.[23] 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

[22] Algebraically, $E/P = (E/D) \times (D/P)$. Our theory predicts one-for-one passthrough of the pure discounting term to $\log(1 + (D/P)) \approx D/P$. A given change in D/P maps to earnings yields according to $\Delta(E/P) = (E/D) \times \Delta(D/P) \approx 2\,\Delta(D/P)$ for a payout ratio of $D/E \approx 0.5$, implying a predicted slope very close to the one observed in the figure.

[23] 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.

Figure 4: 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 1990–2024, or the longest available span for the given country. See Figure 3 for additional details.

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,[24] 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 4 plots 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, $\Delta\widehat{\rho}^*_{t,j}$. In other words, the vertical-axis values are now equal to the difference between the points in the left panel of Figure 3 and the fitted line in that panel. This component 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 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

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

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 1990-2024 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. 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, as can be seen in Figure 2. This period coincides with a large decrease in the U.S.'s net foreign asset position. Japan's estimated ρ^*, meanwhile, increased sharply during this decade, when its net foreign asset position rose. 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 rough cross-country patterns observed in Figure 3. Alternative explanations include relative demographic shifts (e.g., Backus, Cooley, and Henriksen, 2014). We aim to better disentangle these alternative mechanisms empirically in future work.

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.[25]

The starting point for this analysis is a regression for changes in trend real rates, analogous to equation (11) but in differences rather than levels:

$$\Delta r^*_{t,j} = \alpha_0 + \gamma\, \Delta g^*_{t,j} + \beta_j\, \Delta \text{VIX}^2_{t,j} + \Gamma_j + \varepsilon_{t,j}, \tag{13}$$

where Δ denotes a three-year change and where the loading on the VIX term is again country-specific.[26] 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:

$$r^{\text{mkt}}_{t,j} = \alpha_1 + \pi_\rho\, \widehat{\Delta\rho^*_{t,j}} + \pi_g\, \Delta g^*_{t,j} + \pi_V\, \Delta \text{VIX}^2_{t,j} + \Lambda_j + \nu_{t,j}. \tag{14}$$

Table 2 shows the resulting estimates. Before considering our main estimates resulting from (14), we start with a simpler exercise as a benchmark for comparison: we regress three-

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

[26] Coefficient estimates for regression (13) are presented in Table A.1 of Appendix A. 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

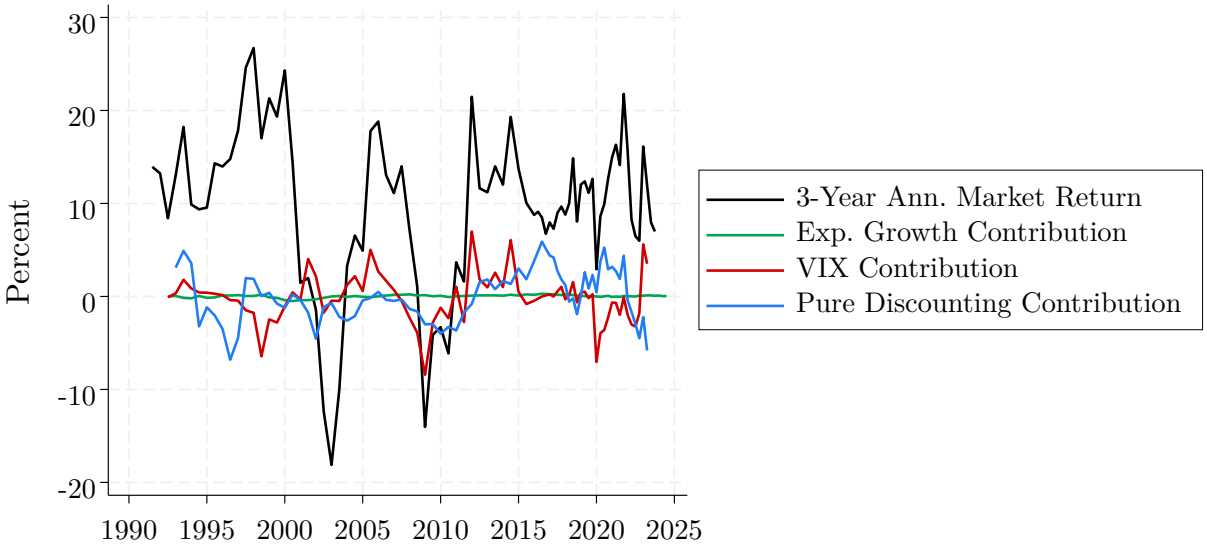
	(1) U.S.	(2) U.S.	(3) All	(4) All
Δ10y yield	4.19		-3.39	
	(3.51)		(2.20)	
Δpure discount ($\widehat{\Delta\rho_t^*}$)		-19.1**		-9.61**
		(7.64)		(3.26)
Δexp. growth		-1.49		16.9*
		(14.0)		(8.82)
$\Delta \text{VIX}^2 \times 100$		-3.08**		-5.44***
		(1.33)		(0.90)
Country FEs	✗	✗	✓	✓
Obs.	74	74	781	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 (13), and the table presents estimates from (14). 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 1990–2024, 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 (13) 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 equation (10), 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 5: 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 (14) 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 (14) is obtained from the first-stage estimation in (13).

the potential for attenuation bias when using the higher-frequency variation in the pure discounting term (see footnote 26 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 5 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 $\widehat{\pi}_V \Delta \text{VIX}^2_{t,j}$ (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 (11)–(12) 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 1990–2024, 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.[27] 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 4.1.

3.4 Cross-Sectional Portfolios

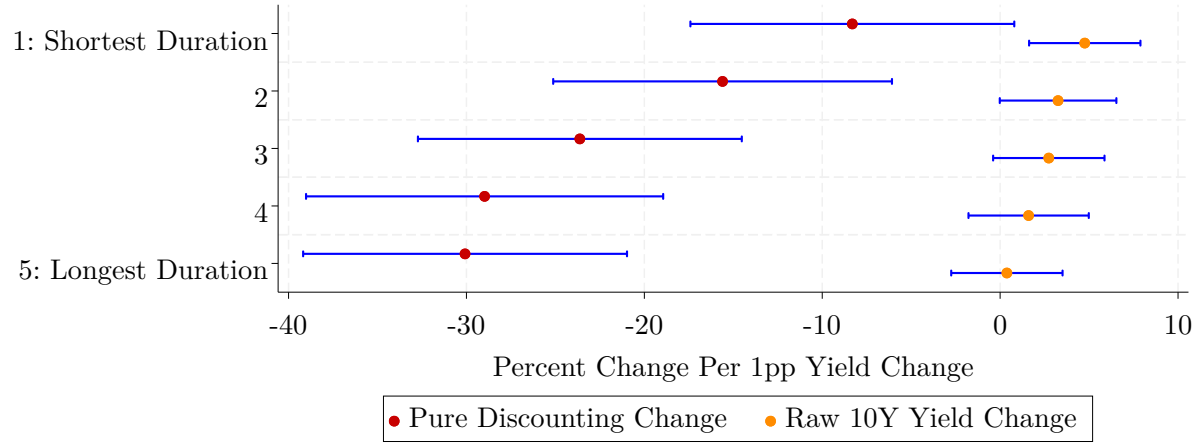
We now turn to the cross-section of stock returns, and study whether firms with different cash flow duration have different exposure to the pure discounting term. A large literature studies the risk and return properties of firms with different cash flow exposure (see Gormsen and Lazarus 2023 and citations 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 litearture.

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.[28]

[27] 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, outperforming well-known predictors from recent literature.

[28] Another standard approach is to proxy for duration by valuation ratios — such as book-to-market ratios — using that a higher valuation ratio is associated with a longer cash flow duration (i.e., "growth" firms are

Figure 6: 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 $r_{t,1} = \alpha_1 + \beta_1 \left(\widehat{\Delta\rho_t^*}\right) + \varepsilon_{t,1}$, 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 (13) using U.S. data. Duration-sorted portfolios and returns are calculated following Gormsen and Lazarus (2023). The sample is 1990–2024.

In Figure 6, 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.[29] 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

long-duration firms). Using book-to-market ratios as the measure of duration does not influence the results presented in this section.

[29]We present corresponding results for the full global sample in Figure A.1.

are not substantially more exposed to raw interest-rate changes than short-duration stocks: 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. Additional 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 4.1), (2) quantify the effect of decreasing interest rates on the value premium (Section 4.2), and (3) help understand the role of an information effect in explaining stock-price responses to monetary policy news (Section 4.3).

4.1 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.[30]

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

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

result, an unadjusted nominal Treasury portfolio may not represent an ideal counterfactual long-term bond return for comparison with equities.[31]

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] very long-run expected growth measures are not in ample supply." Our approach, and our survey data, help disentangle the extent to which bond returns should have passed through to equities, and therefore how puzzling the apparent low duration-matched equity premium should seem.

To do so, 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, the 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

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

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

$$r_{t+1,\rho} = \alpha_t - (n-1)(\rho_{t+1} - \rho_t). \tag{15}$$

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 (10), 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

[31] 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.

$\Delta\rho_t = \rho_{t+1} - \rho_t = 0$. Our approach is to set this value to

$$\alpha_t = \widehat{\rho}_t^* + \mathbb{E}_t[\pi_{t+5}], \tag{16}$$

where $\widehat{\rho}_t^*$ is the pure discounting term estimated from (12), 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).

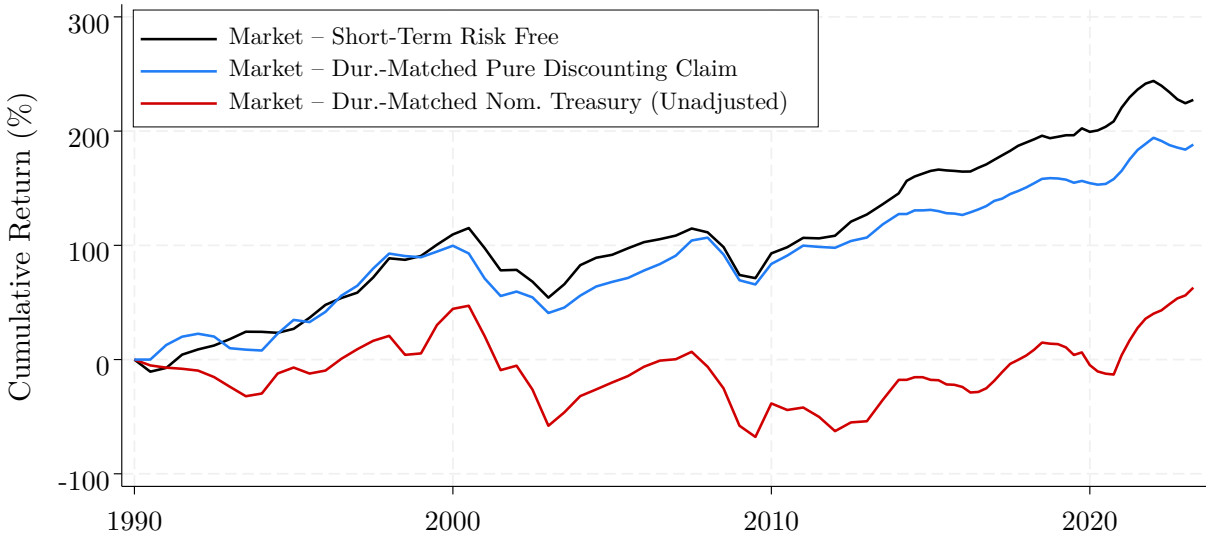
We then calculate the excess return on the market relative to this duration-matched pure discounting claim return $r_{t+1,\rho}$, and cumulate returns over time. We compare this duration-adjusted excess return to both (1) the market return in excess of the short-term nominal risk-free rate, and (2) the market return in excess of a duration-matched nominal Treasury security (unadjusted for changes in growth rates or risk), analogous to van Binsbergen (2024).[32]

These three cumulative returns are plotted in Figure 7. 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 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 7 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

[32] 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 7: 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 (15)–(16) using the estimated $\widehat{\rho}_t^*$ from (12). The red line shows the return relative to an unadjusted nominal Treasury security with duration $\mathcal{D} = 19.1$ years; see footnote 32 for details of construction.

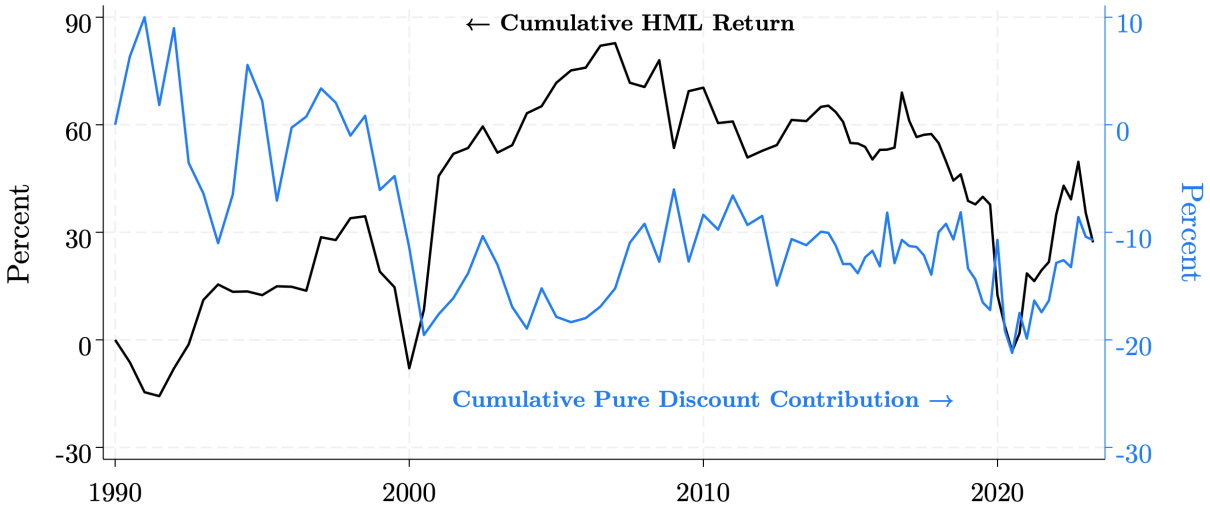
rationalizable: Treasury returns were high in large part because expected growth rates decreased (and, to a lesser extent, because uncertainty increased). The relative performance of stocks over this period is accordingly no longer puzzling once this effect on bond returns is properly accounted for.

4.2 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.[33] 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,

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

Figure 8: 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 (11)–(12) 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 (14)). 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.

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.[34]

We illustrate and quantify the effect of changes in the pure discounting term for the value factor in the U.S. data in Figure 8. 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

[34]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.

equations (11)–(12), controlling for changes in growth rates and uncertainty,[35] 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 sample. In addition, Figure 8 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 9, 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 pure discounting term.[36] 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.

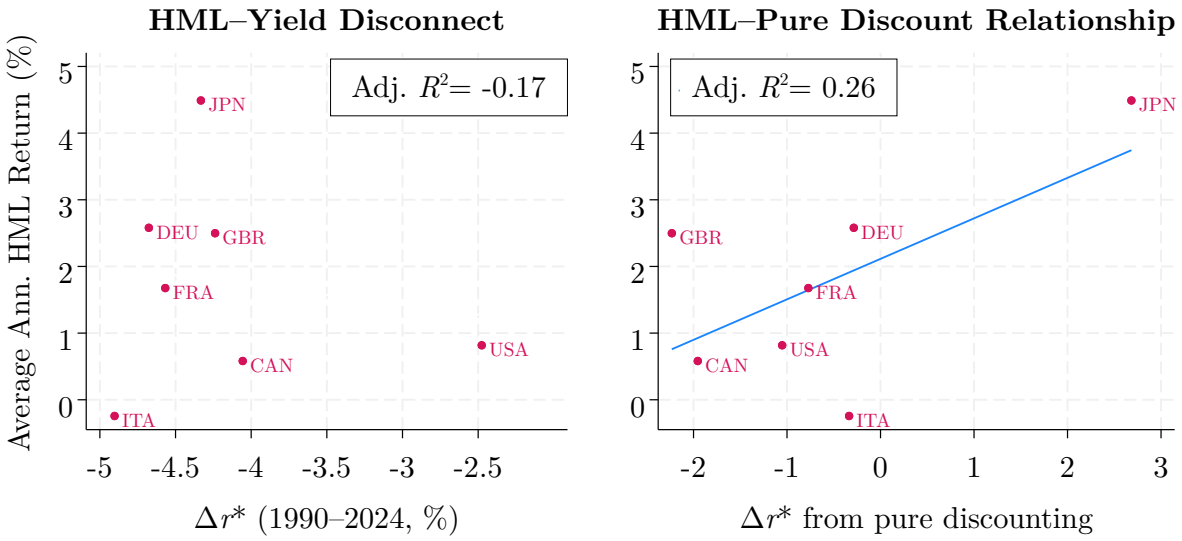
4.3 Unpacking Monetary Policy Shocks

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

[35]This estimation exercise parallels the one in equation (14) for the overall market, with the exception that here we use the change in the residual $\widehat{\rho}^*_{t,j}$ estimated in levels from (11) (rather than the residual from the first-difference estimation in (13)). 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 8. 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 8 would only show 11 equally spaced points).

[36]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.

Figure 9: 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 1990–2024 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 (11)–(12) following the main specification in column (3) of Table 1. The sample is 1990–2024. 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 footnote 20). 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.

if they represent pure discounting shocks (i.e., shocks to ρ_t), 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.[37] 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

[37]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 stock-price responses to interest-rate shocks. Such stock-price responses 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, it turns out, 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).

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 (4). 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.[38] 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 $\Delta y_{t,t+n}$, 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.[39]

[38] Partly as a result of such inertia, our difference-based estimation in Section 3.3 considers three-year changes.

[39] 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

To implement this idea, we start with a slightly modified version of the yield change decomposition in (13). The benchmark yield change we will use in the high-frequency data is a change in the 10-year yield, whereas (13) 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 A.1.[40] Next, we estimate (14) using the three resulting terms from that decomposition. Estimates are again quite similar to the benchmark shown in Table 2.[41]

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.[42] 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 announcements since 1994, a 100-basis-point positive shock mps_t (i.e., a shock scaled so that the impact on the four-quarter Eurodollar futures contract is 100 bps) results in the following:

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

2. 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.

3. An increase in the pure discounting term of $\Delta\rho_t = 29$ basis points. This is again

requires slightly more structure than the one in Nagel and Xu (2024), but it is again similar in spirit.

[40]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).

[41]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.

[42]We thank the authors for posting an updated version of this data set.

significant at 1%, and the regression has a high R^2 of 0.30.

4. 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.[43] 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.

5. 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 theoretically. When implemented empirically using 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

[43] 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.

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.[44] 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, whether in a structural or reduced-form sense, 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.

[44] 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. In addition, in future iterations of this project, we plan to consider the effects of possible forward-looking changes in expected profit shares, as discussed in footnote 12.

Appendix

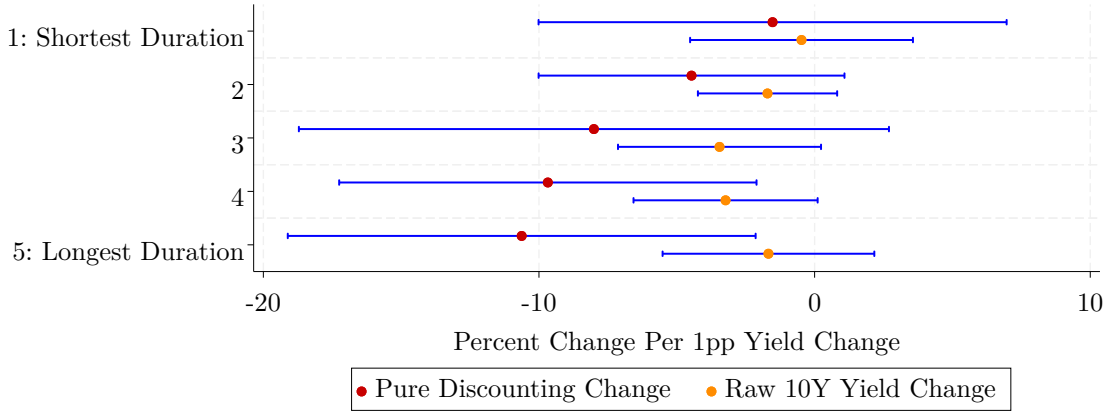
A. Additional Empirical Results

Table A.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 (13), 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 significantly 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 1990–2024.

Figure A.1: Portfolio Exposures to Pure Discount Rate Changes: Global Stocks



Notes: This figure repeats the analysis shown in Figure 6 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 6, with country-level fixed effects. The sample is 1990–2024.

References

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

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*.

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., T. Cooley, and E. Henriksen (2014): "Demography and Low-Frequency Capital Flows," *Journal of International Economics*, 92, S94–S102.

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.

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*.

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., 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*.

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*.

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 (2024): "How the Wealth Was Won: Factor Shares as Market Fundamentals," *Forthcoming, Journal of Political Economy*.

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.

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*.

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 (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.