

Purifying the Equity Premium*

Christopher Polk
London School of Economics
Email: c.polk@lse.ac.uk

Tuomo Vuolteenaho
Email: ptuomov@gmail.com

First Draft: October 2025
This Draft: November 2025

* We thank Jiahong Shi for excellent research assistance and John Campbell, Randy Cohen, Josh Coval, Jussi Keppo, Tom Knox, Sam Thompson, and seminar participants at the London School of Economics for helpful comments.

Purifying the Equity Premium

Abstract

The equity premium has conventionally been defined as the return on stocks minus the return on bills. We decompose this conventional definition into two components: the *pure equity premium*, defined as stocks minus duration-matched inflation-indexed bonds, and the *real term premium*, defined as duration-matched inflation-indexed bonds minus bills. Empirically, we find that the pure equity premium has no “puzzling” features, as it not only has a relatively low Sharpe ratio, but it also has high correlation with measured consumption growth. However, challenges remain: The real term premium is much more volatile than stock returns, and realizations of the pure equity premium and real term premium are very negatively correlated.

I. Introduction

Most investors, researchers, and educators agree that investors seek to earn an expected return premium in excess of the risk-free rate when investing in equities. Logue's (2002) CFA Digest article states that "Investors care about the expected excess return they receive for taking on the risk of investing in equities. This excess return, otherwise known as the risk premium, is the difference between the expected return on common stocks and the risk-free rate of return."

The above quote, however, leaves open the question of what exactly is the relevant risk-free asset with which to compare to equities. It is common in both academia and industry alike to define the equity premium as the return differential between stocks and a short-term nominally risk-free asset. In this paper, we argue that this conventional definition of equity premium is less useful than it could be, propose an alternative to it, and decompose the conventional definition into two economically interpretable components.

The shortcomings of the conventional equity-premium definition stem from two sources, both related to the conventional definition of the risk-free asset. First, the risk-free asset commonly used to compute the realized equity premium, such as the three-month Treasury-bill or LIBOR rate, has a short maturity. In contrast, the stock market is a long-term asset that in principle never matures. Second, it is generally accepted that equities are claims to real assets that produce cash flows that grow over long periods at rates that are stable in real and not nominal terms. Consequently, we argue that the risk-free asset compared to equities should be an inflation-indexed bond, not a nominal bill.

For very short maturity risk-free assets the distinction between real and nominal risk-free cash flows is minor. This is simply because inflation is relatively stable in the short run and usually easy to forecast over short horizons such as three months. However, as the maturity of the risk-free asset increases, the wedge between real and nominal risk-free cash flows grows in importance. Thus, comparisons between equities and long-term nominal bonds become smeared by the long-term impact of inflation. If the duration of the risk-free asset is increased to correspond to that of equities, usually thought to be in

multiple decades, then it becomes increasingly important that the risk-free asset has cash flows that are risk free in real, not in nominal terms.

We decompose the conventional definition of the equity premium into its two components, the *pure equity premium* and the *real term premium*. We define the pure equity premium as the stock market return minus the return on a real risk-free asset that matches the stock market's duration as closely as possible. Specifically, we use the longest-duration inflation-indexed “real” bond available, such as a U.S. Treasury inflation-protected security (TIPS) or a U.K. index-linked gilt (Linkers), and lever that bond up or down, when necessary, to match the duration of the stock market. The real term premium is the return on our long-term real risk-free asset, with duration matching that of equities, funded with the short-term nominal risk-free asset commonly used to compute the equity premium. Thus, the conventional definition of the equity premium is simply the sum of our pure equity premium and real term premium.

By our definitions, the pure equity premium is the return on stocks in excess of an inflation-indexed bond levered to the same duration as stocks. This definition attempts to remove the premium earned by holding long-term assets from the premium earned by holding assets with risky cash flows. We hope that this novel approach allows us to separate the true risk premium due to dividend variability from the return earned simply because of exposure to the longer end of the real yield curve that is inherent in all long-term assets.

We empirically examine the pure equity premium and real term premium in the U.S. and U.K. markets and reconcile them with the conventional definition of the equity premium. Including the U.K. data, in addition to the U.S. data, in our analysis has at least four benefits. First, and most obviously, this increases our effective sample size, as the U.S. and U.K. markets are not perfectly correlated. Second, the inflation-indexed bond market in the U.K. is older, thus allowing us to extend the sample to cover the 1990s. Third, the longest maturity inflation-indexed bond in the U.K. happens to have a real-rate duration relatively close to that of U.K. stocks throughout our sample, requiring only modest leverage for the inflation-indexed bond position to match the real-rate duration of stocks. Fourth, one could argue that the overall economic performance of the U.S. has

been better than expected during the last decades. The much more muted economic performance in the U.K. may help balance our overall sample in that regard.

Our four principal findings are easily summarized. First, we use a valuation model to estimate the duration of stocks at each point in time. Based on our estimates, U.S. stocks have averaged a duration of about 55 years and U.K. stocks about 30 years over our full sample periods. These duration estimates are, however, highly variable. The U.S. duration reached 119 years during the period at the turn of the millennium that has since become known in the industry as the “technology-media-telecom bubble”. The U.K. stock duration peaked at about 71 years in 2015. These estimates confirm the expectation that stocks are very long-duration assets.

Our duration estimates are consistent with those by others who compute equity duration from cash-flow valuation models and much longer than estimates obtained by regressing stock returns on bond returns in a simple regression. We believe that our approach and, thus, our estimates, are correct. Later in this paper, we discuss how that simple regression coefficient is divorced from equity duration by an economically significant omitted variable in the regression.

Second, the pure equity premium is not nearly as “puzzling” as the conventional equity premium. By the equity premium puzzle, academic researchers refer to the tendency of the conventional equity premium being too high and consistent to be explained by reasonable risk aversion. The high Sharpe ratio of the conventional equity premium, defined as its mean divided by its standard deviation, is puzzling to economists because investors’ consumption is estimated to be very smooth and not very correlated with conventional equity premium realizations. Only an implausibly high relative risk aversion coefficient would explain why investors are not, on average, much more heavily allocated to stocks than they actually are. Our below empirical work shows that the pure equity premium has not earned puzzlingly high in-sample Sharpe ratios and can therefore be rationalized with a high, but not implausibly high, relative risk aversion coefficient.

In our sample, the conventional equity premium realized an annualized return of 8.4% and a Sharpe ratio of 0.53 in the U.S. and 3.0% and 0.21 in the U.K. In contrast, the pure equity premium realized an annualized return of 3.8% and a Sharpe ratio of 0.10

in the U.S. and 0.8% and 0.04 in the U.K. We find the U.K. estimates particularly remarkable. Taking on dividend risk simply did not pay off at all in the U.K. over that time. We note that this striking result is over the full sample for the country in which deep and liquid inflation-indexed bond markets have existed the longest.

These in-sample Sharpe ratios of the conventional equity premium imply implausibly high relative risk aversion in the double digits in the U.S. and over five in U.K. for the (representative) investor in the Hansen-Jagannathan (1991) analysis, if we assume a plausible but high volatility of quarterly consumption growth of 1.8% for the representative investor. In contrast, relative risk aversion based on a Hansen-Jagannathan bound computed under the same consumption growth volatility assumption but based instead on realizations of the pure equity premium is approximately two in the U.S. and close to zero in the U.K. Most direct experimental measurements of relative risk aversion put it in the one-to-five range.

Although the representative investor's consumption growth is notoriously poorly measured, when we attempt to approximate it with measured aggregate non-durables and services consumption, we find that the pure equity premium has a higher in-sample consumption growth beta than the conventional equity premium. This result further strengthens our empirical finding that the magnitude of the average pure equity premium can potentially be explained with a reasonable risk aversion coefficient. In summary, the famous equity premium puzzle is not present in the pure equity premium component that eliminates the duration mismatch found in previous work.

Third, even after the historically extreme long-term bond bear market that began in late 2021, the real term premium has provided a relatively high average return and Sharpe ratio in our full sample. The real term premium realized an annualized return of 4.6% and a Sharpe ratio of 0.12 in the U.S. and 2.3% and 0.14 in the U.K. These average returns and in-sample estimated Sharpe ratios are higher than for the pure equity premium; in the U.K., the real term premium has a realized Sharpe ratio that is considerable higher than that of the pure equity premium (0.14 vs. 0.04). These point estimates come with the caveat that annualized Sharpe ratios in our sample sizes (341

months for the U.S. and 455 for the U.K. sample) have standard errors in the neighborhood of 0.19 for the U.S. and 0.16 for the U.K.

Because inflation-indexed bond markets are a relatively recent invention in the context of the sample lengths required to estimate Sharpe ratios, most of our reported point estimates are not statistically different from zero at conventional levels of significance. To reduce the statistical uncertainty of our premia estimates, we introduce a control-variates regression approach that is especially useful for estimating the average real term premium. Our control-variates regression asks the question: What would these premia have been if there had been no change in real interest rates over the sample period? Under the plausible assumption that constant-maturity real yields are strongly stationary, our control-variates regression reduces the standard error of our estimates of the average real term premium by about half. These more efficient estimates indicate a 2/3 and 1/3 split of the mean conventional equity premium between the pure equity premium and the real term premium in the U.S. while allocating the entire meager conventional equity premium in the U.K. to the pure equity premium.

Given that inflation-indexed bonds can convincingly be argued to be the safest asset to a long-term investor, it is not easy to come up with reasons why a long-term investor should require a large positive return premium comparable to the pure equity premium to simply lock in a certain future consumption path. Although explaining this feature of the data with an economic model is outside the scope of the paper, speculating about it is not. In our opinion, we can rule out the liquidity premium as a reason for this. From a bond dealer's perspective, TIPS and Linkers may be relatively illiquid compared to, say, on-the-run ten-year Treasury note. However, from a long-term investor's perspective, especially considering the other sorts of investment options popular with those investors (e.g., private equity), these bonds have very low transaction costs and are traded in extremely liquid and deep markets.

We are also skeptical of explanations of the representative investor having a short horizon such that the high short-term return volatility of inflation-indexed bonds results in that type of investor demanding a positive real term premium. For one thing, asset-liability matching is a common practice by investors in these markets, indicating a long-

horizon perspective. Furthermore, real term premium realizations have a (statistically insignificant) negative correlation with our measured consumption growth, making it a (statistically insignificant) hedge, not a risky asset deserving a risk premium.

Fourth, the pure equity premium and real term premium realizations are extremely negatively correlated. The correlation averages -0.91 in the U.S. and -0.65 in the U.K. and is highly statistically significantly different from zero in both cases. Although there are challenges in reliably statistically testing the time-varying range of these correlations, we believe that it is clear that the conditional correlation has never been positive in our U.S. sample period with a high degree of statistical confidence. In fact, for the last twenty years, the U.S. two-year rolling correlation between pure equity premium and real term premium realizations hovers very close to the average -0.91 level.

We believe that this very strong negative correlation between the pure equity premium and the real term premium has at least two implications for future research on capital markets. In terms of the first implication, simple regression estimates of stock market duration obtained from regressing stock returns on (real) bond returns suffer from a very (negatively) correlated omitted variable, the pure equity premium, and are therefore unlikely to produce reliable estimates of duration, interpreted as the average time to when the present value is received. The correlation is so negative that one can simply say goodbye, or perhaps even farewell, to any hope of measuring stock market cash-flow duration from a simple regression of stock returns on bond returns.

As for the second implication, the observation that the pure equity premium and real term premium are so reliably and persistently negatively correlated ought to put some useful constraints on alternative theoretical asset pricing models designed to fit all observed asset-market data. In our opinion, perhaps tautologically, what is needed is a model in which some unexpected events or “shocks” cause the current value of a fixed real consumption stream to decline while simultaneously either increasing future firm cash flows or reducing the expected future pure equity premium. In other words, we would need an economic model of extreme “flight-to-safety” and “flight-from-safety” behavior. We also speculate that such a model would be inconsistent with our (statistically insignificant) empirical finding that the real term premium has an in-sample Sharpe ratio

comparable to that of the pure equity premium. Exactly what is the “safety” of real bonds, anyway, if that safety requires a consistent and large risk premium?

The rest of the paper is organized as follows. The next section reviews a small subset of the related literature that we find particularly relevant to the above four findings. Section III describes the data. The fourth section explains how we measure duration for stocks. Section V analyzes the properties of these two components of the conventional equity premium. The last section concludes.

II. Related Literature

This paper’s objective is closely related to that of van Binsbergen (2025). He compares the returns on equity claims to matched-duration government bond claims and finds that, duration matched, the historical compensation from bearing aggregate dividend volatility in a duration-neutral fashion has been low. Although van Binsbergen mostly analyzes nominal government bonds, the paper’s research questions are similar to those in this paper.

We reject comparisons of stock returns to nominal bonds on prior grounds as the most useful exercise for understanding the pricing of pure dividend risk. Extremely persistent, and thus difficult-to-model, inflation dynamics drive a wedge between real stock returns and real returns on long-term nominal bonds. Campbell, Shiller, and Viceira (2009), for example, find that when inflation is positively autocorrelated, as it appears to be in the data, the uncertainty regarding future inflation increases with the nominal bond holding period, elevating their real return risk over time as a consequence. This fact makes nominal long-term bond return “volatility puzzles” relatively easy to resolve and makes them somewhat divorced from investigations of inherently real asset returns.

Van Binsbergen’s (2025) section 6.1, which compares stock returns to inflation-indexed government bonds in the U.K., is most related to this paper’s analysis. Van Binsbergen finds (in his Table 8 on page 31) that if U.K. stocks have a duration of twenty years or less, U.K. stocks have earned a lower in-sample average return than U.K. inflation-indexed bonds. Considering our view that U.K. stocks have a duration longer than twenty years, van Binsbergen’s negative duration-matched equity premium result is

arguably more extreme than our headline result of a merely small positive in-sample pure equity premium in the U.K.

It is important to note that van Binsbergen’s (2025) U.K. sample analyzed in his section 6.1 ends in December 2021. That sample endpoint is less than two months after the peak of the bond bull market, which “top-ticked” in November 2021. What followed has been a historically unprecedented government bond bear market, which has taken the longest-maturity U.K. linker yield from -2.5% to $+1.5\%$, an approximately four-percentage-point yield swing. An investment in the longest-maturity U.K. linker lost its entire cumulative excess returns earned since the late 1980s in less than four years. Thus, we suspect that van Binsbergen’s empirical returns relating to real bonds would be materially different if he were to update his analysis to include the most recent data.

Our methodological innovation of assuming a stationary real yield process and estimating the real term premium using a control-variate regression stabilizes these real term premium estimates. We also estimate our mean real term premium in a sample that ends in 2021 and show how the simple sample mean is dramatically altered by just a handful of years, while our control-variate term premium estimates remain relatively stable.

In contrast to van Binsbergen’s (2025) often negative average duration-adjusted equity premium results, Gormsen and Lazarus (2025) find “a sizable duration-matched equity premium” (page 1, abstract) of 6.1% per annum in the U.S. 1990–2023 sample. However, Gormsen and Lazarus’s funding portfolio is a “duration-matched ‘pure discounting claim’” (page 38), and we cannot immediately relate it to any traded fixed-income claim. We are, however, fairly certain that it is not the same as our appropriately levered inflation-indexed bond claim that we call the real term premium, and therefore Gormsen and Lazarus’s duration-matched equity premium is not comparable to our pure equity premium.

We concur with Gormsen and Lazarus (2025) that equity duration is high. Gormsen and Lazarus consider twenty years, which they obtain as “...a likely lower bound for the true duration.” Consistent with this bound, we generally measure a duration for U.S. stocks that is longer than twenty years. We also agree with Gormsen and Lazarus

that the regression coefficient from a simple regression of stock returns on bond returns does not accurately reflect stock market duration, which they empirically demonstrate by producing essentially a zero regression coefficient of stock returns on bond returns in the cross-section of countries.

What is the risk-free asset for a long-term investor? The definition of a risk-free asset depends on the investment horizon. Campbell and Viceira (2001) show that for long-horizon investors concerned with real purchasing power, the truly riskless asset is not a short-term nominal bill but a long-maturity inflation-indexed bond or perpetuity. Campbell, Shiller, and Viceira (2009) further emphasize that inflation-indexed bonds provide the only instrument guaranteeing real returns over extended horizons, offering a benchmark for long-term portfolio analysis. Wachter (2003) formalizes this intuition by proving that as risk aversion rises, the optimal portfolio converges to simply holding long-term real bonds, reinforcing their role as the real risk-free asset. Cochrane (2021) revisits this conclusion, confirming that an inflation-indexed perpetuity remains the theoretical long-run riskless benchmark, consistent with the Campbell–Viceira framework.

There exists a long literature defining duration and measuring it for the stock market. The concept of duration originates with Macaulay (1938), who formally defined bond duration as a weighted average maturity of cash flows, linking price sensitivity to interest-rate changes. Hicks (1939) extended this idea in *Value and Capital*, framing duration as an “average period” that captures how future-oriented an asset’s payoffs are. Gordon (1962) developed the constant-growth dividend discount model, embedding duration-like sensitivity in equity valuation. Leibowitz (1986) introduced “total portfolio duration,” relating equities’ rate sensitivity to their empirical comovement with bonds, pioneering regression-based estimates of equity duration. Bernstein (1995) popularized the link between dividend yield and duration, noting the approximate inverse relationship between the two.

Modern research has continued to refine duration measurement. Dechow et al. (2001) and Schröder and Esterer (2012) adapt bond-duration concepts to equities using forecasted cash flows and implied costs of capital. Weber (2018) estimates firm-level cash-

flow durations directly from accounting data, while Mullins (2020) proposes theoretical models validating multiple approaches.

Empirical work further connects duration to rate exposure. Jiang and Sun (2015) and Choi, Richardson and Whitelaw (2022) estimate equity duration via regressions on Treasury yield changes, documenting heterogeneity across leverage and payout policies. Chen (2022) uses Federal Open Market Committee event studies to isolate discount-rate shocks and infer effective durations. Finally, Li and Wang (2023) construct a “valuation duration” for the aggregate market, showing extreme time variation across boom-bust cycles.

Why is the conventional equity premium considered a puzzle in literature? Mehra and Prescott (1985) identify the “equity premium puzzle,” showing that standard representative-agent models with plausible risk aversion cannot reproduce the historically high average return on equities relative to risk-free assets. Hansen and Jagannathan (1991) formalize this tension by deriving their bound: to match observed market Sharpe ratios, any valid stochastic discount factor must exhibit very high volatility relative to aggregate consumption growth volatility.

Subsequent research proposes mechanisms to reconcile theory and data. Rietz (1988) introduces rare, catastrophic “disaster” states that can justify high equity premia and low real rates, an approach later refined by Barro (2006) using international evidence to show quantitative consistency with historical data. Preference-based explanations include Constantinides (1990), who models habit formation to raise effective risk aversion in downturns, and Campbell and Cochrane (1999), who develop a tractable external-habit framework reproducing both the level and dynamics of the premium without excessive consumption volatility. Kocherlakota (1996) surveys these approaches, concluding that despite such innovations, the coexistence of smooth consumption growth and large risk premia remains fundamentally unresolved. Bansal and Yaron (2004), Bansal, Kiku, and Yaron (2012), and Bansal and Shaliastovich (2013) develop a model of long-run risk where persistent dividend and consumption growth changes are negatively correlated with consumption and dividend volatility changes. They test this model and find that it has

some success in explaining the equity premium as well as other features of historical asset returns.

Past research has measured a positive real term premium from TIPS and index-linked gilts once liquidity distortions are accounted for. Andreasen, Christensen, and Riddell (2019) estimate an affine model incorporating an explicit TIPS liquidity factor and show that the residual real-rate component is time-varying and positive on average. D’Amico, Kim, and Wei (2018) use a joint nominal-real no-arbitrage framework that decomposes TIPS yields into expected inflation, inflation risk, liquidity, and a real term premium, documenting substantial liquidity premia early in the market but a persistent positive real risk premium thereafter. Christensen and Spiegel (2017) similarly find a positive mean real term premium in a TIPS-only affine setting, while Haubrich, Pennacchi, and Ritchken (2012) identify time-varying real bond premia distinct from inflation premia in a joint model of nominal and real yields. Earlier contributions such as Gong and Remolona (1996) and Ang, Bekaert, and Wei (2007) also detect significant real-term-structure variation consistent with positive premia across regimes. Grishchenko and Huang (2012), while primarily concerned with inflation risk premia, also disentangle liquidity and real-yield dynamics, showing that non-zero real premia remain after correcting for market frictions. Collectively, these studies tend to conclude that the U.S. and U.K. inflation-indexed bond markets have exhibited a positive compensation for real-rate duration risk once liquidity effects are stripped out.

Flight-to-safety behavior of government bonds has been analyzed by previous research. Ibbotson and Sinquefeld (1976) document differential performance across U.S. stocks, long-term Treasuries, and bills and provide the historical backdrop for episodes when Treasuries outperform in stress. Bansal and Coleman (1996) introduce a “moneyness” or convenience yield that lowers expected returns on highly liquid, transactions-friendly assets, which is interpretable as a safe-asset premium. Surveying indexed bonds, Pflueger and Viceira (2011) detail how inflation-indexed Treasuries hedge real risks and test expectations-hypothesis implications, clarifying when bonds act as safety assets.

Fleckenstein, Lustig, and Longstaff (2014) identify the TIPS–Treasury “puzzle,” consistent with a convenience premium for nominal Treasuries during safety flights. In equilibrium, Campbell, Pflueger, and Viceira (2020) show that flight-to-safety behavior generates inverse comovement between bond and equity risk premia when bonds hedge marginal utility. High-frequency tests find that volatility and tail risk drive flows into Treasuries: Moise (2020) links volatility shocks to higher equity premia and bond premia consistent with safety demand, while Rubin and Ruzzi (2020) show equity tail risk predicts Treasury returns and is priced. Finally, Schmid, Valaitis, and Villa (2024) analyze optimal government debt management in the presence of such safety demand, highlighting policy trade-offs when Treasuries serve as the preeminent safe asset.

Statistical inference questions tackled by this paper have been previously addressed in related literature. Lo (2002) provides the canonical treatment of the Sharpe ratio’s small-sample properties under i.i.d. returns, deriving its sampling distribution and offering exact and asymptotic tests and confidence intervals and establishing standard inference tools for performance evaluation. Amano (2011) develops a generalized control-variate estimator for dependent (stationary) processes, extending variance-reduction and efficiency gains beyond Monte Carlo settings.

III. Data

We choose the value-weight CRSP market portfolio as our proxy for the U.S. stock market, using Fama and French’s (1993) market factor, as published on Ken French’s website, to measure the return and computing the dividend yield using data from CRSP. We choose the FTSE All-share index as our proxy for the U.K. stock market, taking the return and dividend yield data from WRDS. Our nominal bill return for the U.S. is from Ken French’s website while our nominal bill return is the three-month U.K. Interbank Rate from the Federal Reserve Economic Data (FRED) database. We use Shiller’s 10-year smoothed earnings yield, being careful to undo the interpolation of earnings and use month-end prices. The U.K. 10-year smoothed earnings yield is from Finaeon.

All our primary real bond data, including prices, returns, and yields, come from Bloomberg. We use the longest maturity bond that is outstanding in each country as the

real bond in which we invest at a point in time. We calculate the modified duration of these real bonds, but, because some real bonds have negative yields in our sample, we bound the yield in our modified duration formula at a floor of 1 basis point to simplify computations. Our constant maturity real yield data used in the construction of our control variate comes from the U.S. Treasury and the Bank of England.

We measure U.S. consumption using real per-capita seasonally-adjusted nondurables and services from FRED. Our U.K. consumption is from the U.K. Office of National Statistics and is seasonally-adjusted final real consumption expenditure per head. For each of these series, we compute log consumption growth over quarter t as $\ln((C_{t+1}+C_t)/(C_t+C_{t-1}))$ to align it with calendar-quarter returns.

Table I presents the basic descriptive statistics of our data. Panel A covers the United States for the 2/1997 – 6/2025 period. For the U.S., the average monthly return on TIPS is 0.30%, with a relatively high standard deviation of 3.97%. The aggregate stock market return (RM) exhibits a higher mean of 0.88% and a larger standard deviation of 4.62% per month, consistent with the greater risk and variability typically associated with equities. The nominal risk-free rate (RF) averages 0.18%, showing little variation with a standard deviation of 0.17%, reflecting the relative stability of short-term nominal yields and resulting returns. The 10-year smoothed earnings yield averages 3.26%, with limited dispersion of 0.79%, suggesting fairly stable equity valuations across time. The TIPS yield, at 1.84% on average has a higher standard deviation of 1.18% and ranges from –0.51% to 4.27%. Finally, quarterly log consumption growth averages 0.37%, with a standard deviation of 0.83%. In contrast, previous research has found that the volatility of U.S. stockholders’ quarterly consumption growth is much higher (1.80%, sourced from Table I of Malloy, Moskowitz, and Vissing-Jørgensen (2006)).

Panel B covers the United Kingdom for a longer 8/1987 – 6/2025 period. For the U.K., the mean monthly return on inflation-indexed bonds (LINKERS) is 0.45%, but volatility is higher at 5.46% than in the U.S. TIPS market, suggesting greater sensitivity to macroeconomic shocks, perhaps due to longer duration. The aggregate stock market return (RM) is somewhat lower, at 0.62%, with a similar standard deviation of 4.12%, compared to the U.S. equities. The risk-free rate (RF) averages 0.37%, modestly higher

and more variable than its U.S. counterpart. The 10-year smoothed earnings yield is substantially higher, at 4.96% on average. This lower average U.K. valuation reflects both the earlier sample period and the lower valuation multiples during the common sample period. The LINKERS yield, averaging 1.29%, also displays notable dispersion (standard deviation 1.99%) and includes negative observations. Finally, quarterly log consumption growth in the U.K. averages 0.37%, which matches the U.S. mean. In contrast to the U.S., the U.K. consumption is more volatile at 1.64%.

Figure 1 plots the cumulative excess returns to the U.S. (Panel A) and U.K. (Panel B) stocks and inflation-indexed bonds. Two features stand out in the graphs. First, U.S. stocks went on a ten-year run compared to TIPS in late 2015, during which stocks strongly outperformed TIPS. Second, the Linkers performance in the U.K. was stronger than that of stocks until late 2021, but then the cumulative excess returns on Linkers was wiped out by a crash.

Analogously to Figure 1, Figure 2 plots the yields on U.S. and U.K. stocks and inflation-indexed bonds. The measure of yield for the inflation-indexed bonds corresponds to the real interest rate that the bonds will earn over their life. For stocks, we use the smoothed earnings yield as the yield measure. The prominent features of the yield graphs are the following. Until the year 2000, stock and inflation-indexed bonds were close to each other. Then, from year 2000 to 2021, a relentless trend of declining inflation-indexed bond yields created a large wedge between stock and bond yields. This trend then sharply reversed in late 2021, with inflation-indexed yields skyrocketing and closing the entire gap between stock and inflation-indexed bond yields in the U.S. and at least half of it in the U.K.

IV. Duration of Stocks

The duration of a stock, like that of a bond, is the weighted average maturity of its cash flows, weighted by the present values of those cash flows. Since the cash flow from a stock is its dividend, the duration is the present-value-weighted time at which dividends are received. Under the constant-growth assumption of the Gordon (1962) growth model, the duration of a stock equals its price-dividend ratio. This observation stems from the fact

that the modified duration in the Gordon growth model is the reciprocal of the real growth rate and real discount rate, which in turn equals the price-dividend ratio with the relevant dividend being the next year's dividend. This relation between the dividend yield and equity duration has been known at least since Bernstein (1995) published his book.

This concept of duration treats expected dividends as given and only contemplates a change in the discount rate. Furthermore, the real interest rate is only one component of the discount rate. As such, a stock's duration does not equal the regression beta of a stock's return on real interest rate changes if shocks to expected dividends or risk premia are correlated with shocks to real interest rates. We argue that duration is nevertheless a very useful concept in understanding how far in the future the cash flows from stocks are.

Operating within the present-value model's context, there is an interesting question as to what extent net share repurchases should be included in the concept of dividends in the duration formula. A dynamic trading strategy can have a cash-flow or "dividend" pattern that has an arbitrary duration. For example, one can simply reinvest the dividends and alter the "duration" of the cash flows from this portfolio strategy. It is clear that such trading strategies do not have any causal impact on the duration of the company's cash flows, even if they might produce either large near-term cash flows from stock sales or large long-term cash flows from reinvestment of dividends into additional stock purchases. Whether we trade a single share of stock or not, the underlying company cash-flow duration will remain the same, provided that our trade does not alter the market price of the stock. It is therefore necessary to clearly define what dynamic trading strategy creates the economically most meaningful correspondence with the underlying cash flows of the companies constituting the stock market.

We can easily think of two such trading strategies. The first is simply holding stocks and not participating in any voluntary corporate actions such as open-market share buybacks or seasoned equity offerings. This strategy leads to a duration measure that is the inverse of the dividend yield, i.e., the price-dividend ratio. This is the measure we adopt in this paper.

The second is participating in equity issuance and repurchases pro rata with all of the other owners, which leads to a duration measure that is the inverse of the total payout

yield (i.e., the dividend yield computed including the net share repurchases in the dividend). Although we do not present results corresponding to this dynamic trading strategy, we believe that our main subsequent results are not sensitive to the choice between including vs. not including the net share repurchases in the price-dividend ratio.

In future work, one might consider other economically motivated duration measures. One such measure is the duration of equity cash flows, whether those are paid out of the company or held in cash or marketable securities by the company. Following this chain of reasoning might lead to a duration measure very similar to that of Schröder and Esterer (2012). Although Schröder and Esterer operate with accrual accounting measures, such an equity cash-flow duration measure would likely resemble a cash-based version of their measure, which is the price-to-book ratio divided by the product of the payout ratio and the (accounting) return on equity. For individual companies, such methodological variations might alter the results meaningfully, as there is a wide variation in dividend and share repurchase policies across individual firms. On prior grounds, however, we believe that for the entire stock market such methodological variations would not materially impact our main results.

Another question is how to time the measurement of dividends. Perhaps the cleanest measure of dividends would be the current one-year dividend strip price for the market. Since this price is not easily available for our full sample, we use a version of trailing dividends in our duration formulas. Since we use the measured stock market duration only to lever the inflation-indexed bond to the same duration as the stock market, the exact dividend estimation methodology is unlikely to be consequential. In contrast, when duration is used to predict the stock market, as Li and Wang (2023) do, the exact dividend measurement methodology can be consequential to results.

One could argue in favor of an alternative measure of stock market duration that is based on a time-series regression. Specifically, as mentioned above, one could attempt to measure “duration” as the regression coefficient of stock excess return on the change in real interest rate. This argument and measure fail on two counts. First, based on the present value formula, the value of the stock market depends not only on the real interest rate but also on its risk premium and expected dividends. Thus, controls for shocks to the

risk premium and expected dividends would be needed in that regression. Second, we argue that in order to correctly compute the pure risk premium, one needs to take a stand on the duration of the stock market, leading to a circularity that may be difficult to resolve without resorting to a duration measure computed from the present-value formula. Because of these reasons, in this paper we measure duration using the present value (and some assumptions) and then subsequently measure the correlations between various quantities of interest with duration-matched inflation-indexed bond portfolio.

Table II reports the summary statistics of duration for stocks and real bonds in the U.K. and U.S. Panel A corresponds to the U.S. 2/1997–6/2025 sample. The duration of the longest-maturity TIPS averages 21.5 years, with a time-series standard deviation of 5.0 years. In contrast, the duration of the U.S. stock market is substantially higher, averaging 54.7 years with a larger standard deviation of 16.3 and reaching as high as 119 years. Under our valuation-based framework, U.S. stocks are far more sensitive to changes in discount rates than is the longest TIPS outstanding during this period (holding everything else constant, of course). Figure 3, Panel A, plots the duration of the U.S. stock market and the longest-maturity TIPS.

The SCALEUS metric reported in Table II, Panel A, which we define as the ratio of stock duration to real bond duration, has a mean of 2.8. On average, the U.S. stock market has had its present value about three times further out in the future than the longest-maturity TIPS. The scale factor also varies considerably (standard deviation of 1.2), highlighting fluctuations in both the maturity of the longest TIPS outstanding as well as the relative valuation of stocks and TIPS.

Panel B of Table II summarizes the duration statistics for the U.K. market. Over the 8/1987–6/2025 period, the longest-maturity U.K. inflation-indexed bond (labeled as LINKERS in the table) had an average modified duration of 33.0 years, with a relatively high standard deviation of 12.9 and a range from 18.2 to 52.7 years. This reflects a wide variation in linkers’ maturities as well as variation in the sensitivity of real bonds to interest rates over time for a given time to maturity, stemming from market yields and coupon rates. In contrast to linkers, the U.K. stock market exhibited a slightly lower

average duration of 29.9 years, with lower variability (standard deviation of 9.3) and a maximum of 70.5 years.

Figure 3, Panel B, shows the duration time series for U.K. stocks and linkers. Interestingly, unlike in the U.S., the average stock duration in the U.K. is very close to the average longest-maturity bond duration. The relatively similar duration of U.K. stocks and real bonds leads the SCALEUK variable to have a mean of 1.0, indicating that, on average, the duration of U.K. equities and longest-maturity linkers was about equal over this period. The ratio fluctuated between 0.4 and 2.6, with a standard deviation of 0.5, showing some variability but generally hovering near parity. This finding contrasts with that of the U.S., where stocks have had consistently longer duration than even the longest-maturity TIPS.

Based on prior empirical research using longer time periods, what kind of valuation-based duration should we expect for the U.S. stock market? Fama and French (2002, Table 1 on page 641) report the average dividend yield in the U.S. as 4.70% for the 1872–2000 period. For reference, the corresponding real return on short-term nominal commercial paper over the same period was 3.24%. The duration corresponding to this average dividend yield is 21.3 years. This duration is lower than our 2/1999–6/2025 average of 54.7 years, which is, of course, in the context of, on average, higher stock prices. Thus, we judge our valuation-based duration measures for the U.S. market to be plausible in light of previous research but higher than in the pre-TIPS era.

V. Decomposing the Equity Premium: Pure Equity vs. a Real Term Premium

We define the conventional equity premium as the stock market return in local currency minus the nominally risk-free local currency short-term interest rate. We define the real term premium as the product of two variables. The first is the excess return on inflation-indexed bonds (TIPS in the U.S., Linkers in the U.K.) in excess of the nominally risk-free local currency short-term interest rate. The second is the scale parameter for the country in question as of the beginning of the return measurement period. What this construction achieves is a levered position in inflation-indexed bonds that matches the duration of the corresponding country’s stock market at each point in time. In other words, the levered

position in real bonds now has the duration depicted by the dashed red lines in Figure 3. The pure equity premium is the difference between the conventional equity premium and the real term premium. By construction, the real term premium and pure equity premium sum up to the conventional equity premium each month.

A natural concern arises regarding the assumptions underlying our procedure for constructing a duration-matched real bond portfolio. Specifically, one might question how a long-duration inflation-indexed bond, when levered up, can be treated as a proxy for an unlevered ultra-long bond. After all, leverage itself does not alter the timing of bond cash flows. For the levered bond's return to approximate that of an unobserved ultra-long-duration real bond, it must be the case that yields along the unobserved far end of the real yield curve behave similarly to those at the observed maturities. In effect, we are assuming that the unobserved ultra-long real yields move in parallel with long-term real yields. The validity of our approach therefore hinges on the empirical plausibility of this parallel-shift assumption and on the degree to which the long end of the real yield curve is approximately flat.

By coincidence, in our data the sample-mean SCALEUK parameter equals roughly one, meaning that the leverage adjustment for the U.K. real term premium represents only a small fine-tuning around a naturally matching equity duration. The situation differs in the United States, where the stock market's duration is systematically longer than that of the longest-maturity TIPS and thus SCALEUS systematically above one. Consequently, a more material leverage adjustment is required. As noted above, this procedure implicitly assumes that the real yield curve shifts in parallel across maturities. Empirically, this assumption appears quite reasonable in the data: movements in the 20-year and 30-year constant-maturity TIPS yield series are almost perfectly parallel, suggesting that the unobserved ultra-long segment of the real yield curve would also shift similarly. Figure 4 shows the time series of those two TIPS yields. Focusing on the 300 basis point move after December 2021, one must squint to tell the 20-year and 30-year TIPS yields apart! Hence, the leverage-based duration matching is well grounded in observed yield-curve behavior.

Table III shows the annualized mean, standard deviation, and Sharpe ratio for the conventional equity premium, the real term premium, and the pure equity premium. In the U.S. from 2/1997 to 6/2025, the conventional equity premium averaged 8.44% per annum. In this sample, both the real term premium, 4.60%, and pure equity premium, 3.84%, roughly speaking, contributed equally to the conventional equity premium. Both the real term premium and the pure equity premium were considerably more volatile than the conventional equity premium. The real term premium realized an annualized volatility of 37.71%, while the pure equity premium realized 39.13% in the U.S. sample, which are both more than twice that of the conventional equity premium. This higher volatility of the real term premium and pure equity premium dilutes their Sharpe ratios further compared to that of the conventional equity premium. The conventional equity premium realized an annualized Sharpe ratio of 0.53, while the real term premium realized 0.12 and the pure equity premium 0.10, with only the first one being statistically significantly different from zero.

The second panel of the table shows the corresponding statistics for the U.K. and a longer time period (8/1987–6/2025). The conventional equity premium averaged 3.02% per annum in the U.K., which is less than half of the U.S. point estimate discussed above. In this sample, most of the U.K. conventional equity premium was due to the real term premium, which averaged 2.27%. In contrast to the U.S., the pure equity premium was only approximately a quarter of the conventional equity premium at 0.75%. In contrast to the U.S., where the real term premium and pure equity premium are multiple times as volatile as the conventional equity premium, in the U.K. the series have more equal annualized volatilities: the conventional equity premium 14.29%, real term premium 16.10%, and pure equity premium 19.60%. All of the U.K. Sharpe ratios are much lower than that of the U.S. conventional equity premium: 0.21, 0.14, and 0.04, correspondingly.

Table IV Panel A shows the time-series correlations of the three series for the U.S., and Panel B for the U.K. In the U.S., the real term premium has a correlation of 0.12 and the pure equity premium has a correlation of 0.29 with the conventional equity premium. Interestingly, the real term premium has -0.91 correlation with the pure equity premium, and this negative correlation is highly statistically significantly different from zero. In the U.K., the correlations of the real term premium (0.17) and pure equity premium (0.59)

with the conventional equity premium are higher than in the U.S. The correlation of the real term premium with the pure equity premium is -0.70, which is also highly statistically significantly different from zero.

Figure 5 plots the rolling two-year correlations of the real term premium and pure equity premium estimated from daily data. The averages of the lines are not far from the monthly-data correlations presented in Table IV. The two standard error bounds are also plotted on the graphs. Figure 5 drives home the point that the negative correlation is highly statistically significant, as the higher standard error bound is never close to zero but always deep in negative territory.

The magnitude of the average conventional equity premium is often considered a “puzzle” because of two reasons. First, the conventional equity premium has a high historical Sharpe ratio, especially in the U.S. stock market data. Second, the conventional equity premium is not very correlated with the measured consumption growth. Table V uses quarterly data to estimate the consumption betas of the three series in both countries. We exercise caution in interpreting the results that use U.K. consumption data for two reasons. First, the U.K. consumption data corresponds to a broader set of goods and services than the U.S. consumption data, which includes only nondurable goods and services. Second, as documented by French and Poterba (1991), U.K. investors exhibit much less home bias than U.S. investors, making the often sector-imbalanced U.K. stock index returns less relevant to the U.K. consumer.

Table V Panel A reports U.S. consumption betas. The conventional equity premium has a consumption beta of 2.89, but when the real term premium is subtracted from the conventional equity premium, the resulting consumption beta of the pure equity premium increases to 4.76. The difference is due to the negative consumption beta of the real term premium at -1.71, which however is not statistically significantly different from zero. The U.K. consumption beta patterns are similar to those in the U.S., the main difference being that the less volatile premium series and more volatile consumption growth series in the U.K. produce somewhat lower magnitudes for consumption betas. Motivated by the results of Daniel and Marshall (1997), Table V also reports the consumption betas estimated using one-year and two-year horizon consumption growth

and returns. The results generally get somewhat stronger with longer horizons in the U.S. while remaining mostly inconclusive on the question how the conventional equity premium's consumption beta is split between real term premium and pure equity premium in the U.K.

Another way to judge the consumption beta of equities is to examine the periods during which consumption suffered a large decline or a “crash.” Table VI and Figure 6 concern such periods. Our sample period experienced one large consumption crash, the global financial crisis of the winter of 2008–2009, and one smaller consumption crash, the COVID lockdown crisis in spring 2020. For example, U.K. consumption growth in the first quarter of 2009 was -19%, a serious decline for a long-term risk-averse, consumption-smoothing investor. The U.K. stock market pure equity premium lost approximately -30% in the nine-month period centered around that consumption crash quarter (Figure 6), indicating how risky equities are for a consumption-smoothing investor.

To summarize the results so far, the pure equity premium is an investment in the stock market funded using a real risk-free asset that matches the stock market's duration. The pure equity premium has a lower mean, higher volatility, lower Sharpe ratio, and higher consumption beta than the conventional equity premium that includes a volatile term premium component. It stands to reason that the pure equity premium is less of a puzzle than the conventional equity premium. We turn to Hansen-Jagannathan (1991) bounds to judge how much less of a puzzle it is.

We implement the Hansen-Jagannathan (1991) bounds using the simplest possible consumption-based asset pricing model: additive separable power utility of consumption with constant relative risk aversion and time-preference parameters. According to this model, the minimum coefficient of relative risk aversion that is consistent with the data is the (maximum) Sharpe ratio divided by consumption growth volatility. Table VII shows variations on this bound in our sample. Panel A uses the U.S. data and Panel B the U.K. data. The first column reports the quarterly Sharpe ratios for the series in question, and the following columns show variations on the relative risk aversion bound.

The first bound estimate, γ_1 , assumes 1.8% quarterly log consumption growth volatility. This assumption is based on Table I of Malloy, Moskowitz, and Vissing-

Jørgensen (2006).¹ Malloy, Moskowitz, and Vissing-Jørgensen measure the consumption of people actually participating in the stock market, arguably the most relevant measure of consumption volatility for our calculation. Assuming 1.8% consumption growth volatility, the Hansen–Jagannathan bound of the minimum relative risk aversion coefficient required to explain the observed in-sample Sharpe ratio is approximately fourteen for the conventional equity premium but only approximately two for the pure equity premium in the U.S. sample. In the U.K. sample, with its considerably lower in-sample mean premia, the bound computed from the conventional equity premium is approximately six but close to zero for the pure equity premium.

The second bound estimate in Table VII, γ_2 , uses consumption growth measured from aggregate data. Although we believe that the consumption growth of stockholders is more relevant, we present these bounds based on aggregate consumption growth for completeness. The minimum relative risk aversion coefficient required to explain the U.S. historical pure equity premium is approximately 4.5, reflecting the lower aggregate consumption growth volatility for the whole economy compared to just stockholders.

The third column, γ_3 , is presented as a reference to illustrate the impact of measured consumption beta on the required coefficient of relative risk aversion. This version of the “bound” is an equality relation that adjusts for the measured correlation between the premium and consumption growth realizations. We are simply solving the relative risk aversion coefficient as the reported Sharpe ratio divided by the measured aggregate consumption growth volatility times the measured correlation of aggregate consumption growth with the corresponding premium. This calibration exercise results in a thoroughly implausible relative risk aversion of 110 with the U.S. conventional equity premium and 35 with the U.K. conventional equity premium. In contrast, the pure equity premium has a calibrated relative risk aversion of 22 in the U.S. data and about one in the U.K. data. Although a relative risk aversion of over twenty is still implausible, the required relative risk aversion dropping to one-fifth when the real term premium is

¹ This estimate is found in their working paper, but not the published version. We have confirmed with the authors that the working paper estimate is appropriate to use.

removed from the return stream is nevertheless a large reduction in the magnitude of the equity premium puzzle.

After concluding that the in-sample pure equity premium statistics are much less puzzling than those of the conventional equity premium, we now turn to the real term premium. In a sense, our work above didn't make the conventional equity premium puzzle disappear but simply identified that the puzzle resides not in the pure equity premium but in the real term premium.

Van Binsbergen's (2025) and this paper find that the sample mean real term premium is large, in van Binsbergen's case larger than the conventional equity premium. We find it untenable that the true real term premium would be so large. Both in-sample evidence and economic logic suggest that the real term premium is the safe asset for a long-term investor, with significant diversification benefits even for a short-term investor. The in-sample consumption beta of the real term premium in Table V is (statistically insignificantly) negative both in the U.S. and the U.K. Despite this, the in-sample average term premium is high and comparable in magnitude to the in-sample average pure equity premium. In fact, we can't calibrate a positive relative risk aversion coefficient in Table VII's third column for the in-sample real term premium.

A thing to note is the surprisingly high statistical uncertainty about the mean real term premium. One way to see this statistical uncertainty is simply to recognize the large standard errors in Table III. Another way to see the impact of this statistical uncertainty is to consider the subsample analysis in Table VIII. The first column of the table reproduces the premium statistics from Table III. The second column truncates the sample by excluding the last three and a half years of the sample, and the third column shows the statistics for just the last three and a half years. Looking back at the end of 2021 and examining the longest sample period available then for inflation-indexed bonds, it looked like the real term premium was larger than the conventional equity premium and that the pure equity premium was negative in both countries. This period is approximately the same as what was analyzed by van Binsbergen (2025), and his results are similarly extreme to our subsample results.

The extreme realizations in the third column of Table VIII show how statistically imprecise the pre-2022 results are. The bond market experienced a very significant bear market starting in November 2021, with the real term premium experiencing -45% per annum losses in the U.S. and -21% per annum losses in the U.K. for three and a half years. Without frequent rebalancing, an investor invested in our (levered) real term premium series would have been liquidated by their broker. The addition of a mere three and a half years of additional data to the sample that ends in 2021 moves the average real term premium from 11.54% to 4.60% in the U.S. and from 4.68% to 2.27% in the U.K. Correspondingly, the sample average pure equity premium moves from -2.77% to 3.84% in the U.S. and from -1.73% to 0.75% in the U.K.

Given that we cannot extend the sample further into the past beyond the launch of inflation-indexed bond markets, it is necessary for us to find or develop an estimator of the population mean that is significantly more statistically efficient than the sample mean. This is necessary because an annualized Sharpe ratio of 0.2 would require approximately a century of data for the sample mean to have power to distinguish it from zero (using a standard error of 0.1 as the power requirement). Unfortunately, inflation-indexed bond markets are relatively young, and we only have a quarter or a third of the required century of data, depending on the country.

Since the estimation of the population mean is not exactly a new problem but instead something that others have already taken a crack at, and since various theorems exist about the statistical efficiency of the sample mean, it is clear that we will not improve on the statistical efficiency of the sample mean by an order of magnitude without some additional economic assumptions. In our case, we choose to assume that the constant maturity real yield is a stationary process and thus the constant maturity real yield change has a population mean of zero. This assumption allows us to use a control variates statistical technique which in turn enables a large variance reduction in mean premia estimates, especially for the real term premium.

The specific control variable series we use is the change in 10-year inflation-indexed constant maturity yield. We need to use a constant maturity real yield series instead of a specific bond's yield series because just following a single bond yield change would bake

in some of the term premium and thus wouldn't have a zero mean. The economic interpretation of the control variates regression in this context is the following: On average, constant maturity real yields neither increase nor decrease by assumption. In any given finite sample, those yields may trend either up or down by chance. The control variate regression asks the question what the premia would have been if the interest rates had not trended in either direction by chance.

We re-examine all three premia to make sure that the decomposition of conventional equity premium into the real term premium and pure equity premium is still identically satisfied. We run the following regressions for all three of our premium variables:

$$\text{Realized premium (t)} = \text{mean premium} + b * \text{control variate (t)} + e(t)$$

Taking unconditional expectations of both sides of the regression and assuming that the coefficients are constant and the control variate has a zero mean results in the regression intercept being equal to the mean premium. The variance reduction from this control variates technique is largest when the control variate regression has a high R^2 and the control variate does not exhibit negative autocorrelation and thus has a volatile sample mean. Correspondingly, this control variates technique does not add statistical efficiency if either the R^2 is low or the control variate mean reverts quickly (such that the sample mean control variate has a low variance). Fortunately for our case, realized risk premia are highly correlated with changes in constant maturity real yields and constant maturity real yields are very persistent, providing a good use case for the control variates method. Unreported Monte Carlo simulation experiments confirm this statistical intuition about the power advantage of the control variates technique.

As noted above, one of the assumptions of the control-variates regression is that the regression coefficient on the control variate is constant. In this case, our control variate is the ten-year constant maturity yield change. To ensure a constant regression coefficient, we construct slightly different constant-duration version of our conventional equity premium, real term premium, and pure equity premium series by leveraging and

deleveraging the conventional equity premium duration to 50 years for the U.S. and to 30 years for the U.K. The first row of Table IX Panels A and B shows the sample means for the premia constructed following our original recipe and the second row for the premia constructed using the constant duration recipe. For the U.S., the mean premia are similar; for the U.K., switching to constant duration versions reduces the average real term premium and increases the average pure equity premium.

The regression intercept point estimate and its standard error in Table IX correspond to the conventional equity premium, the pure equity premium, and real term premium mean and standard error of that mean in each of the three regressions. The main premium results in Table IX are on the “Intercept” line under the “Control Variates Regression” heading. For the U.S., the control variates estimate of the mean conventional equity premium for the U.S. is 11.51% (3.05%) per year (standard error in parentheses). The control variates estimates of the mean real term premium and pure equity premium are 4.12% (3.47%) and 7.39% (4.70%). The standard error of mean is cut to about half by the control variates method for the latter two, which is an improvement, but considerable statistical uncertainty about mean premia remains. In the U.K., the control variates premium estimates are 3.01% (2.43%) for the conventional equity premium, -0.44% (1.75%) for the real term premium, and 3.45% (2.87%) for the pure equity premium.

The control variates point estimates are not too different from the sample mean point estimates because, over the full sample, the change in constant-maturity 10-year real yield is not too far from zero. To showcase our application of the control variates technique to mean premia estimation, we also repeat the procedure in a sample that excludes the most recent three and half years and the large bond bear market. Reassuringly, the control-variates intercept estimates are very similar between the full sample and truncated sample. In the U.S., the real term premium control variates estimate from the full sample is 4.12% while that from the truncated sample is 5.51%. The corresponding control variates estimates for the pure equity premium are 7.39% and 6.90%. In contrast, the simple sample means are 4.52% from the full sample and 12.21% from the truncated sample for the real term premium and 7.02% and 0.34% for the pure equity premium, a dramatic change in magnitudes from adding just the last three and half years of data.

VI. Conclusion

This paper decomposes the conventional equity premium, defined as the return on equities minus the return on short-term nominally risk-free bills, into two distinct and economically interpretable components. The first component is the pure equity premium, measured relative to duration-matched inflation-indexed bonds. The second is the real term premium, measured as the return on those long-term real bonds relative to short-term nominal bills. This decomposition separates cash-flow risk from duration risk and allows separate evaluation of those two excess return sources.

Our empirical results from the United States and the United Kingdom yield several notable insights. First, the equity premium puzzle, as conventionally defined, is largely a puzzle about the properties of the real term premium, not the pure equity premium. The pure equity premium in both the U.S. and U.K. samples exhibits modest premia and reasonable in-sample Sharpe ratios. Furthermore, the pure equity premium is positively correlated with measured economy-wide consumption growth. Taking these findings together, the pure equity premium is not a puzzle that demands extreme investor risk aversion.

Second, the real term premium itself is surprisingly large and volatile. In our sample, real term premium volatility is higher than that of the conventional equity premium. Despite representing compensation for locking in real purchasing power over long horizons, most of our average real term premium estimates are positive (although not statistically significantly so).

Third, we find that the pure equity premium and real term premium are strongly and persistently negatively correlated, with a correlation of -0.91 in the U.S. and -0.65 in the U.K. This empirical relationship constrains asset pricing models that attempt to jointly explain stock and bond returns: any plausible model must account for the persistent and highly significant inverse relationship between pure equity cash-flow risk and real term premium realizations.

We also propose a novel application of the control-variates technique for estimating the average risk premium from relatively short samples. Economically plausible true Sharpe ratios, unfortunately, require long return time series for statistically reliable

measurement. Because inflation-indexed bonds are a relatively new instrument, our sample period is necessarily short. To improve the statistical precision of our estimates, we develop a novel application of the control-variates regression to estimate mean real term premia under the assumption that constant-maturity real yields are stationary. Our best estimates of the mean real term premium are 4.12% for the U.S. and -0.44% for the U.K. and of the mean pure equity premium 7.39% for the U.S. and 3.45% for the U.K.

Based on our empirical results, we believe that future theoretical research should aim to build models that can simultaneously generate (i) moderate and consumption-correlated pure equity premia, (ii) volatile real term premia, and (iii) a strong negative correlation between the two. To be consistent with our historical sample, such models would need to capture episodes of flight-to-safety and flight-from-safety dynamics that jointly affect equity and real bond valuations without producing a large negative real term premium.

What economic models might produce the patterns that we observe in the data? Bansal and Yaron (2004) model both a persistent predictable component of consumption and dividend growth rates and the volatility of consumption growth rates. They argue that these key long-run risks, when priced by a representative investor with Epstein-Zin-Weil preferences, solve the equity premium puzzle. Our results can help further discipline that model as not only must one be able to actually detect predictable and persistent variation in consumption growth; an empirically valid version of that model must also produce a mean real term premium that is not too low to be consistent with our and earlier empirical results by Beeler and Campbell (2012).

Gormsen and Lazarus (2025) provide an avenue for future research that may shed light on the extreme negative correlation of our pure equity premium and real term premium realizations. Gormsen and Lazarus decompose the real bond return into three components due to the change in the pure time preference of investors, change in expected very-long-term economic growth, and change in precautionary savings demand. Future research could decompose our pure equity premium into Campbell's (1991) cash-flow news and expected return news and our real term premium into Gormsen and Lazarus's three

components to measure which of the six covariances drives the extreme negative correlation between our pure equity premium and real term premium realizations.

References

- Amano, T. and Taniguchi, M., 2011, Control variate method for stationary processes, *Journal of Econometrics* 165, 20-29.
- Andreasen, M. M., Christensen, J. H. E., and Riddell, S., 2019, The TIPS Liquidity Premium. Federal Reserve Bank of San Francisco Working Paper.
- Ang, A., Bekaert, G., and Wei, M., 2007, The Term Structure of Real Rates and Expected Inflation. National Bureau of Economic Research Working Paper No. 12930.
- Bansal, R., and Coleman, W. J., 1996, A monetary explanation of the equity premium, term premium, and risk-free rate, *Journal of Political Economy* 104, 113-142.
- Bansal, R., Kiku, D., and Yaron, A., 2012, An empirical evaluation of the long-run risks model for asset prices, *Critical Finance Review* 1, 183-221.
- Bansal, R., and Shaliastovich, I., 2013, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies* 26, 1-33.
- Bansal, R., and Yaron, A., 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481-1509.
- Barro, R. J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823-866.
- Beeler, J. and Campbell, J. Y., 2012, The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment, *Critical Finance Review* 1, 141-182.
- Bernstein, R., 1995, Style investing: Unique insight into equity management. New York, NY: John Wiley and Sons.
- Binsbergen, van, J. H., 2025, Duration-Based Stock Valuation: Reassessing Stock Market Performance and Volatility, *Journal of Financial Economics*, forthcoming.
- Campbell, J. Y., 1991, A variance decomposition for stock returns, *The Economic Journal* 101, 157-179.
- Campbell, J. Y., Cochrane, J. H., 1999, By force of habit, *Journal of Political Economy* 107, 205-251.
- Campbell, J. Y., Pflueger, C. E., and Viceira, L. M., 2020, Macroeconomic drivers of bond and equity risks. Working Paper, Harvard University.
- Campbell, J. Y., Shiller, R. J., and Viceira, L. M., 2009, Understanding inflation-indexed bond markets. Brookings Papers on Economic Activity 2009, 79-138.

- Campbell, J. Y., and Viceira, L. M., 2001, Who should buy long-term bonds?, *American Economic Review* 91, 99–127.
- Chen, Q., 2022, Inferring stock duration around FOMC surprises, *Journal of Financial and Quantitative Analysis* 57, 1879–1908.
- Choi, J. J., Richardson, S. A., and Whitelaw, R. F., 2022, Capital structure priority effects in durations, stock–bond comovements, and corporate bond returns. *Review of Asset Pricing Studies* 12, 1–46.
- Christensen, J. H. E., and Spiegel, M. M., 2017, A New Normal for Interest Rates? Evidence from Inflation-Indexed Debt. Federal Reserve Bank of San Francisco Working Paper No. 2017-07.
- Cochrane, J. H., 2021, Portfolios for long-term investors. NBER Working Paper No. 29926.
- Constantinides, G. M., 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.
- Daniel, K. and Marshall, D., 1997, Equity-premium and risk-free-rate puzzles at long horizons, *Macroeconomic Dynamics* 1, 452–484.
- D’Amico, S., Kim, D. H., and Wei, M., 2018, Tips from TIPS: The Informational Content of TIPS Prices, *Journal of Financial and Quantitative Analysis* 53, 395–436.
- Dechow, P. M., Hutton, A. P., Meulbroek, L., and Sloan, R. G., 2001, Implied equity duration: Evidence from analysts’ earnings forecasts. Unpublished manuscript, University of Michigan.
- Fama, E. F., and French, K. R., 2002, The Equity Premium, *Journal of Finance* 57, 637–659.
- Fleckenstein, M., Lustig, H., and Longstaff, F. A., 2014, The TIPS-Treasury bond puzzle, *Journal of Finance* 69, 2151–2197.
- French, K. R. and Poterba, J. M., 1991, Investor Diversification and International Equity Markets, *American Economic Review* 81, 222–226.
- Gong, F., and Remolona, E., 1996, Inflation Risk in the U.S. Yield Curve: The Usefulness of Indexed Bonds, Federal Reserve Bank of New York Research Paper No. 9631.
- Gordon, M. J., 1962, The investment, financing, and valuation of the corporation. Homewood, IL: R. D. Irwin.
- Gormsen, N. J., and Lazarus, E., 2025, Interest rates and equity valuations, working paper.

- Grishchenko, O. V., and Huang, J., 2012, Inflation Risk Premium: Evidence from the TIPS Market, Federal Reserve Board Finance and Economics Discussion Series 2012-06.
- Hansen, L. P., and Jagannathan, R., 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.
- Haubrich, J. G., Pennacchi, G. G., and Ritchken, P., 2012, Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps, *Review of Financial Studies* 25, 1588–1629.
- Hicks, J. R., 1939, Value and capital: An inquiry into some fundamental principles of economic theory. Oxford: Clarendon Press.
- Ibbotson, R. G., and Sinquefeld, R. A., 1976, Stocks, bonds, bills, and inflation: Year-by-year historical returns (1926–1974), *Journal of Business* 49, 11–47.
- Jiang, W., and Sun, J., 2015, Equity duration: A puzzle on high dividend stocks. Working Paper, Columbia Business School.
- Kocherlakota, N. R., 1996, The equity premium: It’s still a puzzle. *Journal of Economic Literature* 34, 42–71.
- Leibowitz, M. L., 1986, Total portfolio duration: A new perspective on asset allocation, *Financial Analysts Journal* 42, 18–29.
- Li, W., and Wang, X., 2023, Valuation duration of the stock market. *Finance Research Letters* 55, 104598.
- Lo, A. W., 2002, The statistics of Sharpe ratios, *Financial Analysts Journal* 58, 36–52.
- Logue, A. C., 2002, The equity premium [Summary of the article by E. F. Fama and K. R. French], *CFA Digest*, 32, 62.
- Macaulay, F. R., 1938, Some theoretical problems suggested by the movements of interest rates, bond yields, and stock prices in the United States since 1856. New York: NBER.
- Malloy, C. J., Moskowitz, T. J., and Vissing-Jørgensen, A., 2006, Long-run stockholder consumption risk and asset returns. Harvard Business School Working Paper No. 08-060, Harvard Business School.
- Mehra, R., and Prescott, E. C., 1985, The equity premium: A puzzle, *Journal of Monetary Economics*, 15, 145–161.
- Moise, A., 2020, Flights to safety and volatility pricing, *Journal of Empirical Finance* 57, 120–138.

- Mullins, G., 2020, Equity duration, *Journal of Portfolio Management*, 46, 27–40.
- Pflueger, C. E., and Viceira, L. M., 2011, Inflation-indexed bonds and the expectations hypothesis, *Annual Review of Financial Economics* 3, 139–158.
- Rietz, T. A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.
- Rubin, A. P., and Ruzzi, H., 2020, Equity tail risk in the Treasury bond market. *Journal of Financial Economics*, 138, 510–536.
- Schmid, L., Valaitis, V., and Villa, A. T., 2024, Optimal management of government debt, *Journal of Monetary Economics*, 137, 123–147.
- Schröder, M., and Esterer, I., 2012, A new measure of equity duration: The duration-based explanation of the value premium revisited, *Financial Markets and Portfolio Management* 26, 123–149.
- Wachter, J. A., 2003, Risk aversion and allocation to long-term bonds, *Journal of Economic Theory*, 112, 325–333.

Table I. Summary Statistics

We report summary statistics for key variables in our analysis. We measure the mean, standard deviation, minimum, and maximum of the monthly return of the on-the-run inflation-indexed bond (TIPS or LINKERS), the monthly return on the aggregate stock market (RM), the monthly nominal risk-free return (RF), the smoothed earnings yield of the aggregate stock market (S&P 500 or FTSE All-Share), the yield of the on-the-run inflation-indexed bond (TIPS or LINKERS), and quarterly log consumption growth. The equity and real bond sample consists of monthly observations from February 1997 to June 2025 for the U.S. (Panel A) and August 1987 to June 2025 for the U.K. (Panel B); the yields are beginning-of-period values. The quarterly consumption growth sample covers 1997Q2 to 2025Q1 for the U.S. (Panel A) and 1987Q4 to 2025Q1 for the U.K. (Panel B). We report estimates in percentage points.

Panel A: U.S. Summary Statistics (Feb 1997 – Jun 2025)				
	Mean	St. Dev.	Min.	Max.
TIPS	0.30	3.97	-18.93	14.61
RM	0.88	4.62	-17.12	13.58
RF	0.18	0.17	0.00	0.56
10-yr Smoothed Earnings Yield	3.26	0.79	1.95	7.05
TIPS Yield	1.84	1.18	-0.51	4.27
U.S. log consumption growth	0.37	0.83	-5.68	4.32

Panel B: U.K. Summary Statistics (Aug 1987 – Jun 2025)				
	Mean	St. Dev.	Min	Max
LINKERS	0.45	5.46	-21.38	24.88
RM	0.62	4.12	-25.77	13.79
RF	0.37	0.30	0.00	1.19
10-yr Smoothed Earnings Yield	4.96	0.98	2.95	8.10
LINKERS Yield	1.29	1.99	-2.60	4.59
UK log consumption growth	0.37	1.64	-13.73	7.61

Table II. Stock and Real Bond Durations

We report summary statistics for stock and real bond durations. We measure the mean, standard deviation, minimum, and maximum modified duration of the on-the-run inflation-indexed bond and the aggregate stock market as well as the resulting scaling factor (SCALEUS or SCALEUK, the ratio of the stock market's duration to that of the real bond). The sample consists of monthly observations from February 1997 to June 2025 for the U.S. (Panel A) and August 1987 to June 2025 for the U.K. (Panel B).

Panel A: U.S. Duration (Feb 1997 – Jun 2025)				
	Mean	St. Dev.	Min.	Max.
TIPS	21.5	5.0	7.6	29.4
STOCK	54.7	16.3	28.9	119.0
SCALEUS	2.8	1.2	1.2	6.9

Panel B: U.K. Duration (Aug 1987 – Jun 2025)				
	Mean	St. Dev.	Min	Max
LINKERS	33.0	12.9	18.2	52.7
STOCK	29.9	9.3	17.7	70.5
SCALEUK	1.0	0.5	0.4	2.6

Table III. Decomposing the Conventional Equity Premium

We decompose the conventional equity premium into a real term premium and a pure equity premium component. We report the average excess return, volatility, and Sharpe Ratio for the aggregate stock market (the *conventional equity premium*); a levered position in the on-the-run inflation-indexed bond (the *real term premium*), where leverage is chosen at the beginning of each month to match the duration of the aggregate stock market; and the difference between those two (the *pure equity premium*). The sample consists of monthly observations from February 1997 to June 2025 for the U.S. (Panel A) and August 1987 to June 2025 for the U.K. (Panel B). We report annualized estimates, with means and volatilities measured in percentage points and standard errors in parentheses.

Panel A: U.S. Decomposition (February 1997 – June 2025)			
	Mean	Volatility	Sharpe Ratio
Conventional Equity Premium	8.44 (3.01)	16.03	0.53 (0.19)
Real Term Premium	4.60 (7.07)	37.71	0.12 (0.19)
Pure Equity Premium	3.84 (7.34)	39.13	0.10 (0.19)
Panel B: U.K. Decomposition (August 1987 – June 2025)			
	Mean	Volatility	Sharpe Ratio
Conventional Equity Premium	3.02 (2.32)	14.29	0.21 (0.16)
Real Term Premium	2.27 (2.62)	16.10	0.14 (0.16)
Pure Equity Premium	0.75 (3.18)	19.60	0.04 (0.16)

Table IV. Premia Correlations

We report the correlation matrix for the conventional equity premium and its two components, the real term premium and the pure equity premium. The sample consists of monthly observations from February 1997 to June 2025 for the U.S. (Panel A) and August 1987 to June 2025 for the U.K. (Panel B). We estimate standard errors (shown in parentheses) using GMM and Newey-West estimates of the spectral density matrix based on four lags.

Panel A: U.S. Correlations (Feb 1997 – Jun 2025)		
	Conventional Equity Premium	Real Term Premium
Real Term Premium	0.12 (0.09)	
Pure Equity Premium	0.29 (0.09)	-0.91 (0.02)
Panel B: U.K. Correlations (Aug 1987 – Jun 2025)		
	Conventional Equity Premium	Real Term Premium
Real Term Premium	0.17 (0.06)	
Pure Equity Premium	0.59 (0.05)	-0.70 (0.05)

Table V. Consumption Betas

We report estimates of consumption growth betas at the quarterly, annual, and two-year horizons for the conventional equity premium and its two components, the real term premium and the pure equity premium. We define these two components in Table III. Newey-West standard errors are in parentheses; we use a lag parameter of four when measuring quarterly and annual betas and eight when measuring two-year betas. The sample consists of quarterly observations from 1997Q2 to 2025Q1 for the U.S. (Panel A) and 1987Q4 to 2025Q1 for the U.K. (Panel B).

Panel A: U.S. Consumption Risk (1997Q2-2025Q1)			
	Quarterly Consumption Beta	Annual Consumption Beta	Two-year Consumption Beta
Conventional Equity Premium	2.89 (0.66)	3.42 (1.06)	2.75 (1.44)
Real Term Premium	-1.71 (1.09)	-3.07 (1.97)	-5.87 (2.69)
Pure Equity Premium	4.76 (1.37)	7.44 (2.23)	12.87 (2.39)

Panel B: U.K. Consumption Risk (1987Q4-2025Q1)			
	Quarterly Consumption Beta	Annual Consumption Beta	Two-year Consumption Beta
Conventional Equity Premium	0.81 (0.41)	1.29 (0.34)	1.09 (0.45)
Real Term Premium	-0.22 (0.23)	0.06 (0.37)	0.09 (0.68)
Pure Equity Premium	0.99 (0.39)	1.15 (0.45)	1.08 (0.76)

Table VI. Premium Realizations in Consumption Crashes

We report the average monthly excess return on the stock market and its two economic components during consumption crashes. The quarters in Panel A where U.S. consumption crashes are the GFC (2009Q1) and Covid (2020Q1) while the quarters in Panel B where U.K. consumption crashes are the GFC (2008Q3) and Covid (2020Q1). Returns are reported in percentage units.

Panel A: Returns in U.S. Consumption Crashes	
	Realized Return
Conventional Equity Premium	-5.14
Real Term Premium	5.85
Pure Equity Premium	-10.99

Panel B: Returns in U.K. Consumption Crashes	
	Realized Return
Conventional Equity Premium	-6.52
Real Term Premium	-0.59
Pure Equity Premium	-5.93

Table VII. Implied Risk Aversion Based on Hansen-Jagannathan Bounds

We estimate risk aversion implied by Hansen-Jagannathan bounds based on the convention equity premium and its two components, the real term premium and the pure equity premium. We report the quarterly Sharpe Ratio and the implied minimum required relative risk aversion coefficient (γ_1) assuming 1.8% quarterly log consumption growth volatility, the implied minimum required relative risk aversion coefficient (γ_2) based on realized quarterly log consumption growth volatility over the sample in question, and the relative risk aversion coefficient (γ_3) implied by the consumption CAPM. The sample consists of quarterly observations 1997Q2 to 2025Q1 for the U.S. (Panel A) and 1987Q4 to 2025Q1 for the U.K. (Panel B). We estimate standard errors (shown in parentheses) using GMM and Newey-West estimates of the spectral density matrix based on four lags.

Panel A: U.S. Risk Aversion Estimates (1997Q2 –2025Q1)				
	Sharpe Ratio	γ_1	γ_2	γ_3
Conventional Equity Premium	0.25 (0.10)	13.66 (5.74)	29.68 (15.21)	110.09 (75.87)
Real Term Premium	0.05 (0.10)	2.77 (5.45)	6.02 (11.52)	<0
Pure Equity Premium	0.04 (0.10)	2.10 (5.56)	4.56 (12.30)	22.03 (61.68)

Panel B: U.K. Risk Aversion Estimates (1987Q1 –2025Q1)				
	Sharpe Ratio	γ_1	γ_2	γ_3
Conventional Equity Premium	0.10 (0.08)	5.75 (4.54)	6.30 (5.59)	35.27 (45.50)
Real Term Premium	0.08 (0.09)	4.63 (5.22)	5.07 (5.88)	<0
Pure Equity Premium	0.00 (0.09)	0.14 (4.95)	0.15 (5.43)	0.82 (30.14)

Table VIII. Sub-sample Premium Estimates

We report sub-sample estimates of the average excess return, volatility, and Sharpe Ratio for the conventional equity premium and its two components, the real term premium and the pure equity premium. The full sample consists of monthly observations from February 1997 to June 2025 for the U.S. (Panel A) and August 1987 to June 2025 for the U.K. (Panel B). We report annualized estimates, with means and volatilities measured in percentage points.

Panel A: U.S. Estimates			
	Feb 1997 – Jun 2025	Feb 1997 – Dec 2021	Jan 2022 – Dec 2025
Mean			
Conventional Equity Premium	8.44	8.77	6.10
Real Term Premium	4.60	11.54	-44.82
Pure Equity Premium	3.84	-2.77	50.92
Volatility			
Conventional Equity Premium	16.03	15.86	17.37
Real Term Premium	37.71	35.32	49.88
Pure Equity Premium	39.13	39.00	37.75
Sharpe Ratio			
Conventional Equity Premium	0.53	0.55	0.35
Real Term Premium	0.12	0.33	-0.90
Pure Equity Premium	0.10	-0.07	1.35
Panel B: U.K. Estimates			
	Feb 1997 – Jun 2025	Feb 1997 – Dec 2021	Jan 2022 – Dec 2025
Mean			
Conventional Equity Premium	3.02	2.95	3.69
Real Term Premium	2.27	4.68	-21.47
Pure Equity Premium	0.75	-1.73	25.16
Volatility			
Conventional Equity Premium	14.29	14.64	10.54
Real Term Premium	16.10	15.85	17.13
Pure Equity Premium	19.60	19.62	18.50
Sharpe Ratio			
Conventional Equity Premium	0.21	0.20	0.35
Real Term Premium	0.14	0.30	-1.25
Pure Equity Premium	0.04	-0.09	1.38

Table IX. Control-variates Premium Estimates

We report control-variates-based estimates of the conventional equity premium and its two components, the real term premium and the pure equity premium. Panels A (US) and B (UK) measure the average excess return and associated standard error for the aggregate stock market, either unlevered, or levered to a constant duration (50 for the U.S., 30 for the U.K.). Our decomposition subtracts the on-the-run inflation-indexed bond, levered to match the duration of the aggregate stock market. The control-variates-based estimate controls for sensitivity to the change in the log yield of a constant 10-year maturity real bond. The full sample consists of monthly observations from February 2003 to June 2025 for the U.S. and August 1987 to June 2025 for the U.K. We report means and intercepts in annualized percentage points with standard errors in parentheses.

Panel A: U.S. Estimates			
(Feb 2003 – Jun 2025)			
Unlevered	Conventional Equity Premium	Real Term Premium	Pure Equity Premium
Mean	10.63 (3.19)	4.04 (8.22)	6.58 (8.06)
Constant Duration	Conventional Equity Premium	Real Term Premium	Pure Equity Premium
Mean	11.54 (3.09)	4.52 (7.16)	7.02 (7.41)
Control Variates Regression			
Intercept	11.51 (3.05)	4.12 (3.47)	7.39 (4.70)
Coefficient	-3.28 (1.14)	-38.50 (1.30)	35.22 (1.76)
R ²	2.6%	76.6%	59.8%
(Feb 2003 – Jun 2021)			
Unlevered	Conventional Equity Premium	Real Term Premium	Pure Equity Premium
Mean	11.46	13.08	-1.62
St. Err.	(3.37)	(8.29)	(8.70)
Constant Duration	Conventional Equity Premium	Real Term Premium	Pure Equity Premium
Mean	12.55 (3.39)	12.21 (7.49)	0.34 (8.22)
Control Variates Regression			
Intercept	12.41 (3.40)	5.51 (3.90)	6.90 (5.34)
Coefficient	-0.81 (1.39)	-39.45 (1.60)	38.64 (2.18)
R ²	-0.3%	73.0%	58.0%

Panel B: U.K. Estimates			
(Aug 1987 – Jun 2025)			
Unlevered	Conventional Equity Premium	Real Term Premium	Pure Equity Premium
Mean	3.02 (2.32)	2.27 (2.62)	0.75 (3.18)
Constant Duration	Conventional Equity Premium	Real Term Premium	Pure Equity Premium
Mean	3.30 (2.48)	0.47 (2.46)	2.83 (3.10)
<u>Control Variates Regression</u>			
Intercept	3.01 (2.43)	-0.44 (1.75)	3.45 (2.87)
Coefficient	-4.63 (0.96)	-14.64 (0.69)	10.01 (1.14)
R ²	4.6%	49.4%	14.3%
(Aug 1987 – Jun 2021)			
Unlevered	Conventional Equity Premium	Real Term Premium	Pure Equity Premium
Mean	2.95 (2.49)	4.68 (2.70)	-1.73 (3.34)
Constant Duration	Conventional Equity Premium	Real Term Premium	Pure Equity Premium
Mean	3.25 (2.64)	3.23 (2.40)	0.02 (3.15)
<u>Control Variates Regression</u>			
Intercept	2.47 (2.61)	0.18 (1.64)	2.29 (2.88)
Coefficient	-4.02 (1.15)	-15.69 (0.72)	11.66 (1.26)
R ²	2.7%	53.4%	17.1%




Figure 1: Cumulative Returns. Panel A plots the cumulative log excess returns on the on-the-run U.S. inflation-indexed bond (TIPS RF) (black solid line) and the U.S. stock market (USRM RF) (red dashed line) from February 1997 to June 2025. Panel B plots the cumulative log excess returns on the on-the-run U.K. inflation-indexed bond (LINKSRF) (black solid line) and the U.K. stock market (UKRMRF) (red dashed line) from August 1987 to June 2025.




Figure 2: Stock and Real Bond Yields. Panel A plots the yield of the on-the-run U.S. inflation-indexed bond (TIPS) (black solid line) and the earnings yield on the SandP 500 (US STOCK) (red dashed line) from February 1997 to June 2025. Panel B plots the yield on-the-run U.K. inflation-indexed bond (LINKERS) (black solid line) and the earnings yield on the FTSE All-share index (UK STOCK) (red dashed line) from August 1987 to June 2025.

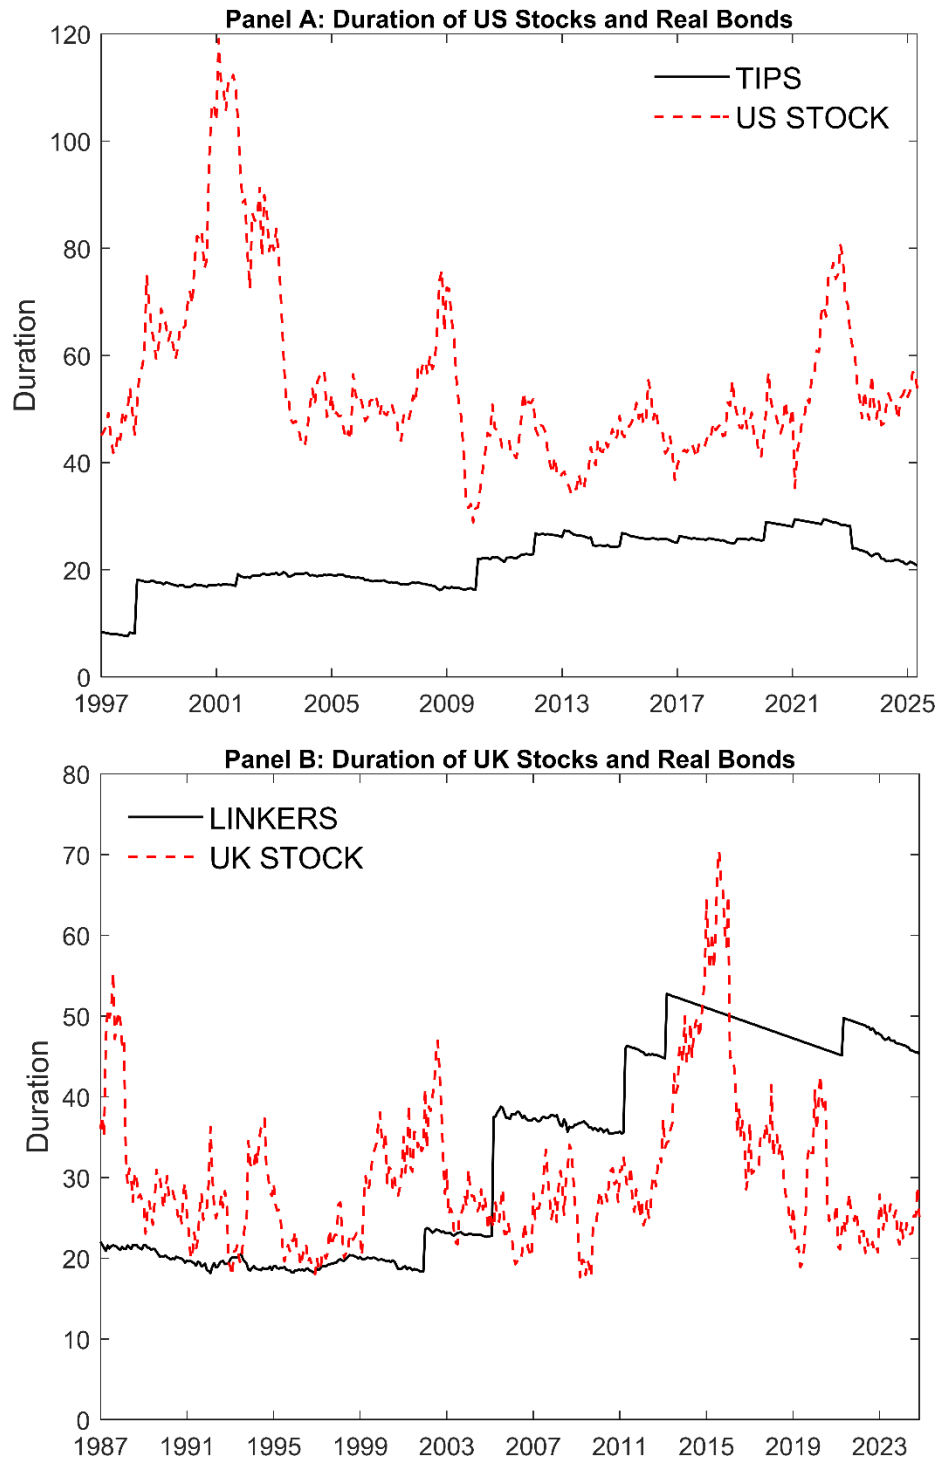


Figure 3: Stock and Real Bond Durations. Panel A plots the modified duration of the on-the-run U.S. inflation-indexed bond (TIPS) (black solid line) and the U.S. aggregate stock market (US STOCK) (red dashed line) from February 1997 to June 2025. Panel B plots the modified duration of the on-the-run U.K. inflation-indexed bond (LINKERS) (black solid line) and the U.K. aggregate stock market (UK STOCK) (red dashed line) from August 1987 to June 2025.




Figure 4: U.S. Constant-maturity Real Bond Yields. We plot the daily yields on constant-maturity 20-year (red dashed line) and 30-year (black solid line) TIPS over the period of our sample where they are available (February 22, 2010 to June 30, 2025).

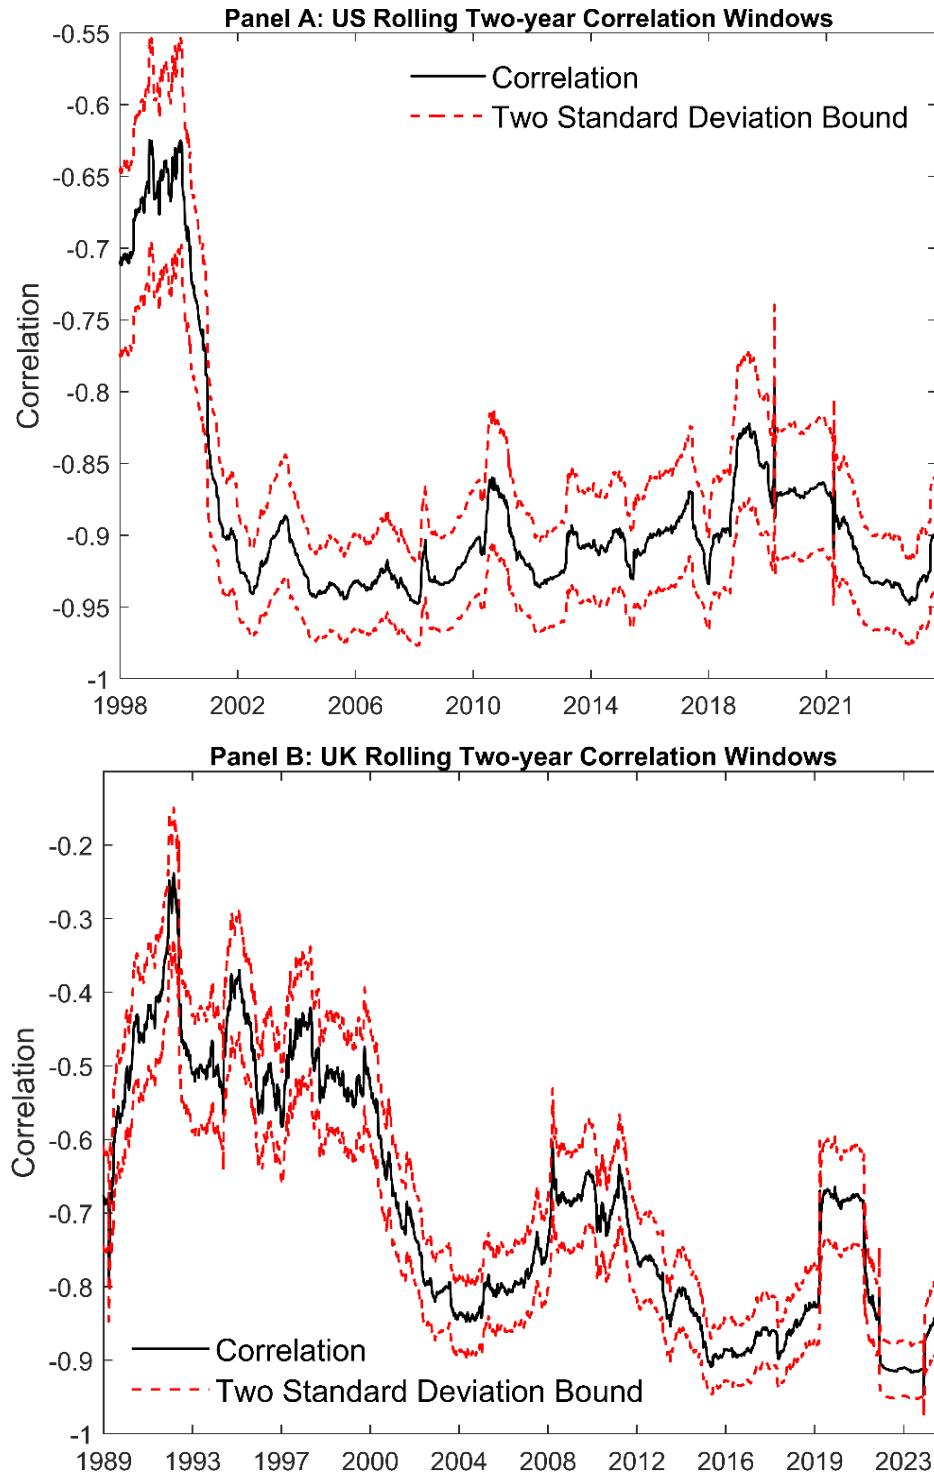


Figure 5: Time-varying Correlations. Panel A plots the daily return correlation between the pure equity premium and the term premium components of the excess return on the U.S. stock market from February 1999 to June 2025. Panel B plots the corresponding correlation for the U.K. market from August 1989 to June 2025. We estimate these correlations using a two-year rolling window with overlapping log-return observations aggregated over three trading days to account for nonsynchronous trading.





Figure 6: Consumption Crashes. We plot the realized returns on the pure equity premium component of the U.S. (Panel A) and U.K. (Panel B) stock market during consumption crashes. During the three quarters centered around Covid and the GFC, U.S. (UK) consumption growth was -7.49% (-18.97%) and -1.36% (-5.69%).