Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

A Stock Return Decomposition Using Observables

Benjamin Knox, Annette Vissing-Jorgensen

2022-014

Please cite this paper as:
Knox, Benjamin, and Annette Vissing-Jorgensen (2022). "A Stock Return Decomposition Using Observables," Finance and Economics Discussion Series 2022-014. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2022.014.

A Stock Return Decomposition Using Observables

Benjamin Knox[1] and Annette Vissing-Jorgensen[2]

[1]Federal Reserve Board

[2]Federal Reserve Board and CEPR

Abstract

We propose a method to decompose stock returns period by period. First, we argue that one can directly estimate expected stock returns from securities available in modern financial markets (using the real yield curve and the Martin (2017) equity risk premium). Second, we derive a return decomposition which is based on stock price elasticities with respect to expected returns and expected dividends. We calculate elasticities from dividend futures. Our decomposition is an alternative to the Campbell-Shiller log-linearization which relies on an assumption about the log-linearization constant (ρ). An application to the COVID crisis in 2020 reveals that risk premium changes drove much of the crash and rebound in the S&P500 while a fall in long-term real yields drove a strong positive return for 2020 as a whole.

This version: March 2, 2022

Emails: ben.knox@frb.gov and Annette.Vissing-Jorgensen@frb.gov. Benjamin Knox acknowledges support from the FRIC Center for Financial Frictions (grant no. DNRF102). We are grateful to Tolga Cenesizoglu (EFA discussant), Lars Lochstoer (WFA discussant) and seminar participants at Copenhagen Business School, UT Austin, UCSD, UIUC, New York University, Johns Hopkins University, Indiana University, Cornell University, Carnegie-Mellon University, Bundesbank and Russian Central Bank/HSE/NES for very useful feedback. We thank Magnus Dahlquist and Markus Ibert for sharing their data on asset manager return expectations. The views in this paper should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System.

I. Introduction

A central theme in asset pricing is what types of news drives realized asset returns. A large literature combines the log-linearization of Campbell and Shiller (1988) with a VAR approach as in Campbell (1991) to decompose stock return variance into components coming from dividend news, expected return news, and the covariance of the two. In this paper we propose a decomposition not of return variance but of the realized return for a given period. This allows for an interpretation of the movement of the stock market period by period. It can also be used to assess potential risks to the market ex-ante, similarly to the way investors use duration to assess risks ex-ante in the bond market.

Our approach relies on two main ideas. First, we make the simple observation that in modern financial markets, a lot of information about the real yield curve and the equity risk premium is observable. The term structure of the real riskless rate can be measured out to around 30 years from nominal Treasury yields and inflation swaps, or from inflation-indexed Treasury (TIPS) yields. Using the nominal rate on interest rate swaps combined with inflation swaps one can observe real yields out as far as 40 years. The term structure of the equity risk premium is not directly observable but Martin (2017) provides a lower bound on the equity risk premium based on S&P500 index options. He argues that this lower bound is approximately tight and thus is close to the actual equity risk premium.[1] While Martin studies the equity risk premium out to 1 year, this can be extended out to around 2 years in recent years, based on available S&P500 options. If fluctuations in the equity risk premium have a large transitory component, fluctuations in the first two years will account for a substantial part of equity risk premium news. We supplement Martin's equity risk premium approach with data for asset managers' equity risk premium estimates out to year $t+10$, obtained from Dahlquist and Ibert (2021). This provides information on less transitory movements in equity risk premium.

Second, to utilize the availability of rich discount rate data, we develop a simple new stock return decomposition (Result 3) which is straightforward to map to available data. Our approach starts from the present value formula that expresses the stock price as a function of expected

[1]We supplement Martin's analysis with theoretical analysis of how the *change* in the Martin lower bound relates to the true change in the equity risk premium. In particular, we show that for the CRRA log-normal case, the same parameters that ensure that the lower bound is in fact a lower bound (Martin's negative correlation condition) also ensure that the change in the lower bound is smaller than the change in the true risk premium. This suggests that our approach will understate the role of risk premium changes for realized returns to the extent that Martin's lower bound is not tight.

dividends and expected returns. Under a set of assumptions, we derive the elasticities at date t of the stock price with respect to the expected return for year $t+k$ and with respect to the expected dividends for year $t+k$. We show that these elasticities are functions only of dividend strip weights (i.e., the fractions of the stock market paid for dividends at various maturities). Using $w_t^{(n)}$ to denote the portfolio weight of the n^{th} dividend strip in the overall stock market, the elasticity of the stock price at t with respect to the expected return in a future year $t+k$ is $-\sum_{n=k}^{\infty} w_t^{(n)}$, i.e., minus the portfolio weight of dividends from $t+k$ onward. This is intuitive, since the present value of dividends to be received before $t+k$ are unaffected by increased discounting in year $t+k$. The elasticity of the stock price at t with respect to the expected dividend at $t+k$ is simply the dividend strip weight for that dividend $w_t^{(n)}$. Intuitively, the effect of a given percentage change in an expected dividend on the stock price depends on how large that expected dividend is prior to the change, measured by its weight in the overall stock market.

To implement our method, we calculate dividend strip weights from dividend futures out to year 10 and show how dividend strip weights can be estimated past year 10 if expected returns and expected growth are approximately constant past year 10. Crucially, we show how market data on dividend futures can be used to estimate the market's perception today of the ratio G/R at the terminal date, where G denotes the gross expected growth rate and R the gross expected return, both past year 10. The dividend futures, combined with discount rate data, also provide information about expected dividends out to 10 years. Armed with elasticities and observable measures of changes to expected returns from real yields or the equity risk premium (and some data on changes to expected dividends), we are able to provide a decomposition of a given daily realized stock return into real yield curve news, equity risk premium news, near-term dividend news, and long-run news capturing news past the measured horizons. Most cash flow news will enter this residual given the modest importance of the first 10 years of dividends for the value of the stock market.

We compare our elasticities approach to the Campbell-Shiller log-linearization which could also be used for a period-by-period stock return decomposition given observable inputs. We argue that our approach maps more directly to data. Our expected return and expected dividend elasticities can be calculated from dividend futures data, while calculation of expected return and expected dividend elasticities based on the log-linearization approach requires an assumption about the value of the log-linearization constant $\rho = \frac{1}{1+D/P}$ (e.g., the last value or the historical average). Because the Campbell-Shiller log-linearization is done around the same value of ρ in all time periods, it implicitly assumes a constant D/P ratio going forward. In our approach, the elasticities out to the horizon of available dividend futures impose no restrictions on the D/P ratio. Furthermore, elasticities past this horizon are calculated based on a market-implied measure of the G/R ratio.

As discussed above, by implementing a stock market decomposition (ours or the CS log-linearization) using observables, one can measure the contributions of various return components for the realized stock return in a particular time period. A standard VAR estimation also produces time series of cash flow news and discount rate news. Could one use a VAR-based analysis to understand stock market developments in a given time period? We think the answer is no, in the sense that this would be less informative. Specifically, a VAR-based return decomposition is typically used to estimate whether stock market fluctuations on average, over a long estimation sample, tend to be driven more by cash flow news or discount rate news. This objective aligns well with the assumptions made. One assumes a time-invariant VAR model structure, with constant regression coefficients. The drawback of using a VAR for a period-by-period decomposition of returns is that VARs interpret all movements in the predictors equally. For example, suppose the price-earnings ratio falls in a given period of interest and that declines in the price-earnings ratios predict higher stock returns over the VAR estimation sample. A VAR approach to understanding the realized return for the period will then suggest that there was negative discount rate component to the realized stock return even if the decline in the price-earnings ratio in this particular period was due mainly to negative news about future cash flows. By contrast, since our approach does not rely on regressions it does not impose any constraints on the mix of discount rate news and cash flow news in a given period.[2] Similarly, because our decomposition is based on observable data and does not require any regressions, our approach allows equity market researchers to decompose realized stock returns in a way that is conceptually similar to that used in event studies of yield changes in bond markets. See, e.g., Krishnamurthy and Vissing-Jorgensen (2011) for a decomposition of yield changes around quantitative easing announcement dates into components.

The question of understanding stock market movements in a given period gained particular interest during the COVID crisis and recovery so we illustrate our approach with an application to 2020. Figure 1 graphs the cumulative return on the S&P500 index over the year 2020.[3] The

[2] Recent work in the VAR literature allows for time-varying coefficients, see e.g. Bianchi (2020).

[3] The figure is based on the total S&P500 return index form Bloomberg (SPXT Index) which index includes dividends. For the year, the dividend yield accounts for 1.8% of the total return.

cumulative S&P500 return was -33.9% from the start of 2020 up to March 23, 2020. This was followed by a sharp recovery and the market ended the year up by 22.1%. The dramatic moves in the stock market led to a series of questions. What drove the sharp decline in the market? Why did the market recover so fast despite the lingering health crisis? Why was the market recovery so strong, with a large gain for the year?

An emerging literature has approached these questions by seeking to measure cash flow news for 2020. Landier and Thesmar (2020) analyze analyst earnings forecasts up to May 2020. They estimate a counterfactual path for the stock market which assumes unchanged discount rates (and payout ratios) and uses dividend expectations updated daily based on updates to analyst forecasts for 2020-2022 earnings. The 2022 earnings are central for their terminal value calculation. They find that dividend news was modest, around -5% by March 23, 2020, and became more negative past March 23, 2020. This is in sharp contrast to the large crash and fast recovery of the actual stock market. Cox, Greenwald, and Ludvigson (2020) study the COVID crisis using the estimated structural asset pricing model of Greenwald, Lettau, and Ludvigson (2019) in which the value of the stock market is expressed as GDP $\times$ [corporate profits/GDP] $\times$ [stock market value/corporate profits]. They also conclude that it is difficult to explain the V-shaped trajectory of the stock market over the COVID crisis with cash flow news. A central argument is that, based on data from the Survey of Professional Forecasters as of May 2020, GDP was expected to fall by about 10% in 2020Q2, but was expected to increase in 2020Q3. The COVID shock to GDP was thus expected to be quite transitory implying that GDP (and thus earnings) changes can explain only a small fraction of the crash and recovery. Gormsen and Koijen (2020) study dividend futures. They show that changes to the value of dividends out to year 10 can account for little of the stock market crash (given their modest weight in the market and the realized decline in dividend futures values) and none of the recovery up to July 20, 2020. They argue that longer-maturity dividends are likely to be only modestly affected by the COVID crisis, implying that changes to their present value and thus to the overall market may have been driven mostly by discount rate news.

We infer from these papers that, to the extent it is possible to measure cash flow news, such news does not appear able to explain much of the stock market decline or recovery in 2020. Is this due to the difficulty of measuring expected cash flows, or can market movements instead be explained by discount rate news? Our decomposition allows us to shed light on this question by taking what is, in essence, the opposite approach of that taken by prior work on 2020. We try to measure discount rate news rather than cash flow news. The focus on discount rates also allows us to provide a more granular decomposition of discount rate effects by distinguishing between (a) the impact of real yield curve news and the impact of equity risk premium news, and (b) the relative importance of short-term and long-term discount rates. Indeed, the application of our decomposition to 2020 reveals four key facts.

First, the equity risk premium increased sharply until March 18 and had a substantial role in the market crash. We estimate that from the start of the year up to March 18, the equity risk premium for the one-year horizon increased from 2.6% to 15.7%, with further increases in the year-2 risk premium. Together, the increase in the risk premium for the first two years contributed -14.3 percentage points to the stock return up to March 18 (which was -26%). Using quarterly data on asset manager equity risk premium estimates out to the 10-year horizon, we estimate that equity premium news can account for almost all of the market decline in 2020Q1. During the recovery period, equity risk premia decline quickly. An "A-shaped" pattern for the equity risk premium thus helps explain the V-shaped pattern of the stock price. Equity premia remain somewhat higher at the end of 2020 than at the start of the year.

Second, with the exception of an upward spike in long rates from March 9-18, real riskless rates drop dramatically across all maturities and do not recover by the end of the year. The 10-year real riskless rate declines over 100 bps over the year and real forward rates fall even out to the 40-year horizon. For the year 2020, the decline in the term structure of real rates out to year 40 contributes a 20.9% increase in the stock market.

Third, changes to expected dividends out to year 10 have a modest effect on the market, contributing minus 2.5% percentage to the stock return over the year and never more than minus 4.5% percent during the year. The small contribution of observed expected dividend changes to overall returns is unsurprising given that the first 10-years of dividends generally contribute only about 20% of the value of the stock market.

Fourth, there is some negative residual (the long-run news) in the crash but almost none for the year as a whole. Since the majority of cash flow news enters the residual in our approach, this finding is consistent with the prior literature finding only a modest role for cash flow news.

The outline of the paper is as follows. Section II derives our stock return decomposition. Section III compares it to the Campbell-Shiller log-linearization. Section IV-VI describes the implementation of our decomposition and applies it to 2020, while Section VII implements it for the financial crisis as well as for the July 2004-December 2020 period of available data. Section VIII concludes.

II. A new stock return decomposition

We derive a simple decomposition of the realized stock market return (over a short period) into real yield curve news, equity premium news and cash flow news. Our approach relies on calculating first derivatives and elasticities of the stock price to expected returns and expected dividends at various maturities. The percentage change in an expected return or exp dividend multiplied by the elasticity of the stock price to this variable is then its contribution to the aggregate stock price movement.

After presenting the new approach in this section, Section III compares the price elasticities from our approach to those one would obtain based on the Campbell-Shiller (CS) log-linearization. We argue that our approach has the benefit that the elasticities are function only of dividend strip weights which can be calculated from dividend futures, whereas the CS log-linearization requires an assumption about the log-linearization parameter ρ.

A. Background definitions

We begin with some definitions of prices, dividends and returns. Unless otherwise noted, all variables are in real terms. Start from the present value formula of the stock market

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

where $P_t^{(n)} = \frac{E_t[D_{t+n}]}{R_{t,n}}$ is the value of the n^{th} dividend strip, i.e., the present value at time t of the expected dividend paid out at time $t+n$ with $R_{t,n}$ the n-period cumulative gross discount rate at time t.

The one-period gross return of the n^{th} dividend strip is given by $R_{t+1}^{(n)} = \frac{P_{t+1}^{(n-1)}}{P_t^{(n)}}$ where $P_{t+1}^{(0)} = D_{t+1}$. The cumulative gross discount rate on the expected dividend paid out at time $t+n$ can then be expressed as

$$R_{t,n} = E_t \prod_{k=1}^{n} R_{t+k}^{(n-k+1)} \tag{2}$$

because $R_{t,n}$ is the time t expected hold-to-maturity return, and the n-period holding return is the product of one-period returns from time $t+1$ to $t+n$.

The one-period gross return on the market can be expressed as the value-weighted average of the one-period gross returns on all dividend strips

$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t} = \sum_{n=1}^{\infty} w_t^{(n)} R_{t+1}^{(n)} \tag{3}$$

where $w_t^{(n)} = \frac{P_t^{(n)}}{P_t}$ is the weight of the n^{th} dividend strip (the present value of the expected dividend paid out at time $t+n$ relative to the overall stock market value).

B. The effect of expected return changes on the stock price

To show the effect of expected returns on the stock price we first make the following assumptions on the returns to dividend strips.

Assumption 1. *The realized returns on a dividend strip are independent across time periods, conditional on information known at date t*

$$E_t \prod_{k=1}^{n} R_{t+k}^{(n-k+1)} = \prod_{k=1}^{n} E_t R_{t+k}^{(n-k+1)}.$$

Assumption 2. *The expected gross return on a dividend strip is proportional to the expected gross return on the market*

$$E_t R_{t+k}^{(n)} = b_t^{(n)} E_t R_{t+k}$$

where $\sum_{n=1}^{\infty} w_t^{(n)} b_t^{(n)} = 1$ *must hold due to equation (3).*[4]

With these assumptions, the effect of expected return changes on stock returns is as follows.

Result 1 (expected return news and stock returns).
Under assumptions 1 and 2, the effect of an instantaneous change to the expected gross return on the market for year $t+k$ *on the stock price can be expressed as:*

$$\frac{\partial P_t / P_t}{\partial E_t R_{t+k}} = -\frac{1}{E_t R_{t+k}} \sum_{n=k}^{\infty} w_t^{(n)}. \tag{4}$$

[4] A simple but more restrictive version of Assumption 2 is that all dividend strips have the same expected gross return for a given period, which is then equal to expected gross return on market. In this case $b_t^{(n)} = 1$ for all t and n such that $E_t R_{t+k}^{(n)} = E_t R_{t+k}$.

Therefore, the elasticity of the stock price with respect to the expected gross stock return in year $t+k$ is

$$\psi_{t,k}^{R} = \frac{\partial P_t / P_t}{\left(\partial E_t R_{t+k}\right) / E_t R_{t+k}} = -\sum_{n=k}^{\infty} w_t^{(n)} \tag{5}$$

with a one percent increase in the expected gross return in period $t+k$ generating a $-\sum_{n=k}^{\infty} w_t^{(n)}$ percent stock return today.

Proof: See appendix.

Result 1 is related to the standard bond pricing formula that links bond price changes to duration and yield curve shifts. However, rather than studying the effects of parallel shifts in the yield curve, we derive the effect of a change to the expected gross return for one future year.

To see the intuition, consider an increase in the expected gross return for period $t+k$, $E_t R_{t+k}$, of one percentage point. With the higher discount rate for year $t+k$, all dividends to be paid at $t+k$ or later will now be discounted by one percentage point more when we discount back from $t+k$ to $t+k-1$. Therefore, if there were no dividends before date $t+k$, then the expected return elasticity of the stock price, $\psi_{t,k}^{R}$, would simply be -1. However, if there are dividends before date $t+k$, their present value is unaffected by the change in the expected return for year $t+k$, resulting in a smaller effect of $E_t R_{t+k}$ on P_t. The factor $\sum_{n=k}^{\infty} w_t^{(n)}$ captures the fraction of today's price P_t that is due to dividends at date $t+k$ and later.

In terms of the assumptions required for Result 1, Assumption 1 states that realized returns on a dividend strip are independent across time periods, conditional on information known at date t. Importantly, this does not rule out time-variation in expected returns and expected returns for different maturities can update in a correlated fashion. What needs to hold is that *realized* returns in one year for a dividend strip are not informative for realized returns in another year on that same dividend strip, conditional on what is known at t. For example, $E_t\left[R_{t+1}^{(2)} R_{t+2}^{(1)}\right] = E_t\left[R_{t+1}^{(2)}\right] E_t\left[R_{t+2}^{(1)}\right] + cov_t\left(R_{t+1}^{(2)}, R_{t+2}^{(1)}\right)$. Thus, the assumption holds for horizon $n=2$ if $cov_t\left(R_{t+1}^{(2)}, R_{t+2}^{(1)}\right) = 0$, i.e., if the distance of $R_{t+1}^{(2)}$ from its conditional mean is uninformative for the distance of $R_{t+2}^{(1)}$ from its conditional mean.

Assumption 2 states that the expected return on a dividend strip is proportional to the expected return on the market. If $b_t^{(n)} = 1$ for all n, then the equity term structure (a plot of the expected annualized return on dividend strips against dividend strip term) is flat at time t and all dividend strips have the same expected return. This expected return is equal to the expected return on the market. However, the assumption is less restrictive than this, and allows for upward sloping ($b_t^{(m)} > b_t^{(n)}$ if $m > n$) or downward sloping ($b_t^{(m)} < b_t^{(n)}$ if $m > n$) term structures of equity returns. The assumptions also allow for time series variation in the term structure (see Gormsen (2021)) as the maturity dependent proportional factors are conditional on t.

As an alternative to assuming proportional gross returns, suppose instead that expected returns on dividend strips followed a CAPM, possibly with non-zero alphas:

$$E_t\left(R_{t+k}^{(n)}\right) - R_{t+k}^f = \alpha_t^{(n)} + \beta_t^{(n)}\left[E_t\left(R_{t+k}\right) - R_{t+k}^f\right]$$

In this case, Result 1 would change to Result 1^{Alt}

$$\frac{\partial P_t/P_t}{\partial E_t R_{t+k}} = -\sum_{n=k}^{\infty}\frac{\beta_t^{(n)}}{E_t R_{t+k}^{(n)}} w_t^{(n)}$$

The main difference from Result 1 is that $\beta_t^{(n)}$ in the numerator may differ substantially from one for some maturities. The net effect of this compared to Result 1 is ambiguous. van Binsbergen, Brandt, and Koijen (2012) and van Binsbergen and Koijen (2017) find that betas are increasing in n (and below one for low n). This has two effects in Result 1^{Alt}. First, since the sum on the right hand side only starts at $n = k$, betas that are increasing in n will imply that the average beta used in the sum is above 1, thus making $\frac{\partial P_t/P_t}{\partial E_t R_{t+k}}$ more negative compared to Result 1. Second, within the sum, betas that are increasing in n will imply that the largest weights $w_t^{(n)}$ are multiplied by the smallest betas, thus making $\frac{\partial P_t/P_t}{\partial E_t R_{t+k}}$ less negative compared to Result 1. Given this ambiguity, and given that Result 1^{Alt} is harder to implement in practice as it requires estimates of $\beta_t^{(n)}$ and $E_t R_{t+k}^{(n)}$, we proceed with Result 1.

C. The effect of expected dividend changes on the stock price

We next characterize the effect of expected dividend changes on stock returns as follows.

Result 2 (dividend news and stock returns).
The effect of an instantaneous change to the expected dividend for year $t+k$ on the stock price can be expressed as:

$$\frac{\partial P_t/P_t}{\partial E_t D_{t+k}} = \frac{1/P_t}{R_{t,n}} = \frac{1}{E_t D_{t+k}} w_t^{(k)}. \tag{6}$$

Therefore, the elasticity of the stock price with respect to the expected dividend at $t + k$ is

$$\psi_{t,k}^{D} = \frac{\partial P_t / P_t}{\left(\partial E_t D_{t+k}\right) / E_t D_{t+k}} = w_t^{(k)} \tag{7}$$

with a one percent change in expected dividend $t + k$ leading to a $w_t^{(k)}$ percent stock return.

Result 2 states that a one percentage point change in a dividend accounts for a percentage change in the stock price that is equal to that dividend's weight in the overall stock market valuation. Since dividends further into the future are discounted more, dividend strip weights $w_t^{(k)}$ will typically be decreasing in maturity k, implying that a one percent change to a later dividend has a smaller effects on today's stock price than a one percent change to an earlier dividend.

D. The new stock return decomposition

Results 1 and 2 show the effect of an instantaneous change in an expected return or an expected dividend on today's price. We implement them using daily data, both because a short time period maps better than a longer one (e.g., weeks or months) to the idea of instantaneous changes, and because we are interested in understanding the evolution of the stock market day by day.

Assuming that, over a one-day period, the realized dividend yield equals the expected dividend yield for that day, we can express the realized return on day d as

$$\text{Realized return}_d = \text{Expected return}_d + \text{Unexpected capital gain}_d.$$

Using Results 1 and 2, we then have the following decomposition:

Result 3 (Realized return decomposition).

$$\begin{aligned}
& \textit{Realized return}_d \\
= \; & \textit{Expected return}_d + \underbrace{\sum_{k=1}^{\infty} \left[-\Sigma_{n=k}^{\infty} w_t^{(n)} \right] \frac{\partial E_t R_{t+k}}{E_t R_{t+k}}}_{\textit{Discount rate news}_d} + \underbrace{\sum_{k=1}^{\infty} w_t^{(k)} \frac{\partial E_t D_{t+k}}{E_t D_{t+k}}}_{\textit{Cash flow news}_d} \\
= \; & \textit{Expected return}_d + \underbrace{\sum_{k=1}^{\infty} \left[-\Sigma_{n=k}^{\infty} w_t^{(n)} \right] \frac{\partial f_{t+k}}{E_t R_{t+k}}}_{\textit{Yield curve news}_d} + \underbrace{\sum_{k=1}^{\infty} \left[-\Sigma_{n=k}^{\infty} w_t^{(n)} \right] \frac{\partial e_{t+k}}{E_t R_{t+k}}}_{\textit{Risk premium news}_d} + \underbrace{\sum_{k=1}^{\infty} w_t^{(k)} \frac{\partial E_t D_{t+k}}{E_t D_{t+k}}}_{\textit{Cash flow news}_d}
\end{aligned} \tag{8}$$

where the second line uses that the expected gross stock return for year $t+k$ can be expressed as

$$E_t R_{t+k} = f_{t+k} + e_{t+k} \tag{9}$$

where f_{t+k} denotes the (gross real) forward rate for a risk-free 1-year investment in year $t+k$ and e_{t+k} denotes the equity risk premium for year $t+k$. Result 1 holds whether changes to $E_t R_{t+k}$ are due to changes in expectations of f_{t+k} or e_{t+k}.

In practice, data for real yields, the equity risk premium and expected dividends is not observed to infinite maturities. This limitation on data availability means that in actual implementations one computes

$$\begin{aligned} & \text{Realized return}_d \qquad (10) \\ = \; & \text{Expected return}_d + \underbrace{\sum_{k=1}^{40}\left[-\Sigma_{n=k}^{\infty} w_t^{(n)}\right]\frac{\partial f_{t+k}}{E_t R_{t+k}}}_{\text{Riskfree rate news}_d} + \underbrace{\sum_{k=1}^{2}\left[-\Sigma_{n=k}^{\infty} w_t^{(n)}\right]\frac{\partial e_{t+k}}{E_t R_{t+k}}}_{\text{Risk premium news}_d} + \underbrace{\sum_{k=1}^{10} w_t^{(k)}\frac{\partial E_t D_{t+k}}{E_t D_{t+k}}}_{\text{Cash flow news}_d} \\ & + \text{Residual}_d \end{aligned}$$

where the residual term reflects the daily returns unexplained by the observable data. In US data, the maximum maturity of available data on real risk-free rates, risk premium and expected dividends are 40, 2 and 10 years respectively (see subsequent sections for the data descriptions). The residual will include discount rate news past the horizons stated as well as the majority of the cash flow news (given the modest role of the first 10 years of dividends for the stock price).

E. Dividend strip weights and stock market elasticities

In implementing Result 3, a central element is the dividend strip weights $w_t^{(n)}$. It is well known that dividend strips (which are not traded) can be valued from dividend futures (e.g. van Binsbergen et al. (2013)). Since dividend futures pay off at maturity $(t+n)$, dividend strips and dividend futures prices are related by

$$P_t^{(n)} = F_{n,t} / \left(1 + y_{t,n}^{\text{nom}}\right)^n \tag{11}$$

where $F_{n,t}$ denotes the date t price of a dividend future paying the nominal dividend for period $t+n$ at $t+n$ and $y_{t,n}^{\text{nom}}$ is the riskless nominal yield at date t for a n-period investment. In this expression, $F_{n,t}$ is nominal (since actual dividend futures contracts pay the nominal dividend) and therefore they are discounted using the nominal yield $y_{t,n}^{\text{nom}}$. Using the dividend futures prices, we can then express the dividend strip weights as

$$w_t^{(n)} = \frac{P_t^{(n)}}{P_t} = \frac{F_{k,t}/\left(1+y_{t,k}^{\text{nom}}\right)^k}{P_t}. \tag{12}$$

This highlights how Result 3 is easily mapped to data when dividend future prices are available.

In practice, dividend futures are traded on the S&P500 index out to a maximum maturity of 10 years. Assumptions are therefore needed to estimate dividend strip weights past this maturity. We assume a Gordon growth model for dividends beyond year $t+10$, with G_t^L denoting the expected annual gross dividend growth rate past year $t+10$ and R_t^L is the expected annual gross return past year $t+10$. Both rates are, as of time t, assumed to be constant in expectation across all periods past year $t+10$. Under this assumption, the value of "long-term dividends", L_t is

$$\begin{aligned} L_t &= \sum_{k=11}^{\infty} \frac{E_t D_{t+k}}{R_{t,k}} \\ &= P_t^{(10)} \left(\frac{G_t^L}{R_t^L} + \left(\frac{G_t^L}{R_t^L}\right)^2 + \dots \right) \\ &= P_t^{(10)} \left(\frac{G_t^L}{R_t^L - G_t^L} \right) \end{aligned} \tag{13}$$

Rearranging in terms of the ratio of the long-run gross dividend growth rate to the long-run gross expected return,

$$\frac{G_t^L}{R_t^L} = \frac{L_t}{L_t + P_t^{(10)}}. \tag{14}$$

The value of long-term dividends, L_t, is observed from the difference between aggregate stock market value and the sum of dividend strip values up to 10 years maturity

$$L_t = P_t - \sum_{k=1}^{10} P_t^{(n)} \tag{15}$$

and $P_t^{(10)}$ is observable as the price of the 10th dividend strip. The right hand side of equation (14) is therefore fully observed. The ratio $\frac{G_t^L}{R_t^L}$ can thus be calculated from observables and we arrive at the following result.

Result 4 (**Long-maturity dividend strip weights estimates**).

Assuming a Gordon growth model for dividends beyond year 10 and using available data on dividend futures, dividend weights beyond 10 years in maturity can be expressed as:

$$w_t^{(k)} = w_t^{(10)} \left(\frac{G_t^L}{R_t^L} \right)^{k-10} \qquad \textit{for } k > 10. \tag{16}$$

where ratio of the long-run growth rate to expected return is calculated from equation (14).

Since $G/R \simeq 1 + G - R = 1 - D/P$ under constant expected dividend growth and expected returns, our assumptions about G^L and R^L are equivalent to making assumptions about the value of D/P at the horizon from which expected dividend growth and expected returns are assumed constant.

To illustrate the relation between dividend strip weights and elasticities, Figure 2 plots the dividend strip weights on the first day of 2020 in the top left hand panel, and the cumulative sum of weights in the top right hand side panel.[5] Dividends up to 10 (30) years were a combined weight of 17% (44%) of the total stock market value on this date. This highlights the large fraction of stock market value that is generated from long-maturity dividends, and thus the long duration of the stock market overall (a point also made in van Binsbergen (2020) and Gonçalves (2021)). As of the start of 2020, the duration of the stock market was 46 years based on the dividend strip weights in Figure 2.

Figure 2 plots the dividend elasticity of the stock price in the bottom left hand panel and the return elasticity of the stock price in the bottom right hand panel. These figures shows how percentage changes in expected dividends and expected returns at various maturities effect the aggregate stock market's instantaneous return. The dividend elasticity of the stock price is the same as the dividend weight in the panel above. The return elasticity of the stock price is $\psi_{t,k}^R = -\sum_{n=k}^{\infty} w_t^{(n)} = -\left(1 - \sum_{n=1}^{k} w_t^{(n-1)}\right) = \sum_{n=1}^{k} w_t^{(n-1)} - 1$ and therefore equals the value in the figure above, minus one. Increases in expected future returns lead to instantaneous negative realized returns.

[5] We obtain the dividend futures prices from Bloomberg, which are available from 2017 and out to a maturity of 10 years. To get constant-maturity prices for years 1 through to 10, we interpolate across the prices of different contracts each day, following the norm in the literature (see van Binsbergen et al. (2012), van Binsbergen et al. (2013), van Binsbergen and Koijen (2017), Gormsen et al. (2021)).

III. Comparison to the Campbell-Shiller log-linearization

As an alternative to our approach that starts from the present value formula, one could have used the Campbell and Shiller (1988) (CS) log-linearization to calculate price elasticities and thus to implement a period-by-period return decomposition with observables. We compare these two approaches before turning to the implementation of our approach.

Let $r_t = \ln R_t$, $p_t = \ln P_t$, and $d_t = \ln D_t$, then it follows

$$
\begin{aligned}
r_{t+1} &= \ln\left(P_{t+1} + D_{t+1}\right) - \ln P_t \\
&= \ln\left(P_{t+1}\left[1 + \frac{D_{t+1}}{P_{t+1}}\right]\right) - \ln P_t \\
&= p_{t+1} - p_t + \ln\left(1 + \exp\left(d_{t+1} - p_{t+1}\right)\right)
\end{aligned}
$$

The CS log-linearization is based on a first-order Taylor approximation of
$\ln\left(1 + \exp\left(d_{t+1} - p_{t+1}\right)\right)$ around a particular value of $d_{t+1} - p_{t+1}$, often the historical average value of $d - p$ denoted by $\overline{d-p}$:

$$
\begin{aligned}
& \ln\left(1 + \exp\left(d_{t+1} - p_{t+1}\right)\right) \\
= \; & \ln\left(1 + \exp\left(\overline{d-p}\right)\right) + \frac{\exp\left(\overline{d-p}\right)}{1 + \exp\left(\overline{d-p}\right)}\left[d_{t+1} - p_{t+1} - \left(\overline{d-p}\right)\right] \\
= \; & k + (1-\rho)\left[d_{t+1} - p_{t+1}\right]
\end{aligned}
$$

with

$$
\begin{aligned}
\rho &= \frac{1}{1 + \exp\left(\overline{d-p}\right)} && (17) \\
k &= -\ln\rho - (1-\rho)\ln\left(\frac{1}{\rho} - 1\right) && (18)
\end{aligned}
$$

This implies,

$$
r_{t+1} = p_{t+1} - p_t + k + (1-\rho)\left[d_{t+1} - p_{t+1}\right] \Longrightarrow p_t = \rho p_{t+1} + k + (1-\rho)\, d_{t+1} - r_{t+1} \qquad (19)
$$

Iterating forward, and assuming the no bubble condition that $\lim_{j\to\infty} \rho^j p_{t+j} = 0$, then

$$p_t = \frac{k}{1-\rho} + \sum_{j\geq 0} \rho^j (1-\rho) d_{t+1+j} - \sum_{j\geq 0} \rho^j r_{t+1+j}. \tag{20}$$

with this expression also holding in expectation. Assuming returns and dividends are log-normally distributed, then

$$\begin{aligned} p_t =& \frac{k}{1-\rho} + \sum_{j\geq 0} \rho^j (1-\rho) \ln E_t D_{t+1+j} - \sum_{j\geq 0} \rho^j \ln E_t R_{t+1+j} \\ &- \frac{1}{2}\left(\sum_{j\geq 0} \rho^j (1-\rho) Var_t\, d_{t+1+j} - \sum_{j\geq 0} \rho^j Var_t\, r_{t+1+j}\right) \end{aligned} \tag{21}$$

since $E(\ln X) = \ln E(X) - \frac{1}{2} Var(\ln X)$ for a log-normal random variable X. Given equation (21), the effect of expected return and expected dividend changes based on the CS log-linearization is as follows.

Result 1^{CS} (expected return news and stock returns based on the CS log-linearization). *Using the Campbell-Shiller log-linearization and assuming returns are log-normally distributed, the elasticity of the stock price with respect to the expected return in period $t+k$ is:*

$$\psi_k^{R,CS} = \frac{\partial P_t / P_t}{(\partial E_t R_{t+k}) / E_t R_{t+k}} = \frac{\partial p_t}{\partial \ln E_t R_{t+k}} = -\rho^{k-1} \tag{22}$$

Result 2^{CS} (dividend news and stock returns based on the CS log-linearization). *Using the Campbell-Shiller log-linearization and assuming dividends are log-normally distributed, the elasticity of the stock price with respect to the expected dividend at date $t+k$ is:*

$$\psi_k^{D,CS} = \frac{\partial P_t / P_t}{(\partial E_t D_{t+k}) / E_t D_{t+k}} = \frac{\partial p_t}{\partial \ln E_t D_{t+k}} = (1-\rho)\rho^{k-1} \tag{23}$$

One can think of Result 1^{CS} and Result 2^{CS} as restricted versions of Result 1 and 2 in the following sense. When iterating (19) forward to derive the price expression, the Campbell-Shiller approach implicitly log-linearizes $\ln(1 + \exp(d_{t+k} - p_{t+k}))$ around the same value of the log dividend-price ratio for all k. This is equivalent to saying that the resulting Campbell-Shiller price expression assumes that, standing at date t, investors expect the dividend-price ratio to be the same at each future date $t+k$. The simplest example of a setting generating a constant dividend-price ratio is one with constant expected growth and constant expected returns. The result below shows that if we were to impose these assumptions on our framework (starting from today's date), it would give identical values for the elasticities of the stock price as the Campbell-Shiller approach if the latter log-linearized around today's D/P ratio.

Result 5 (**Comparing impact of expected return news and dividend news in our approach and Campbell-Shiller log-linearization**).
If (a) returns and dividends are log-normal, (b) expected gross dividend growth and expected gross stock returns are constant going forward (equal to G and R, respectively), and (c) the CS log-linearization is done around the log of today's dividend-price ratio, then our approach simplifies to the CS approach:

$$\rho = \frac{G}{R} \tag{24}$$

$$w_t^{(n)} = (1-\rho)\,\rho^{n-1} \tag{25}$$

$$\psi_k^R = -\sum_{n=k}^{\infty} w_t^{(n)} = -\rho^{k-1} = \psi_k^{R,CS} \tag{26}$$

$$\psi_k^D = w_t^{(k)} = (1-\rho)\,\rho^{k=1} = \psi_k^{D,CS} \tag{27}$$

Proof: See Appendix.

This result makes it clear that the CS log-linearization imposes additional assumptions compared to our approach. From a practical perspective, the main difference is that rather than having to make an assumption about the most accurate value of ρ to use in the CS log-linearization, the dividend strip weights in our approach are observable from dividend futures out to year 10 and, as shown in Result 4, one can use the observable values of the stock price at t and of short-term dividends at t to back out the ratio G_t^L/R_t^L needed for the calculation of dividend strip weights past year 10.

Conceptually, the most accurate value of ρ to use in the CS log-linearization would be based on an estimate of the typical dividend-price ratio going forward, but is this best proxied by a historical average, the latest value, or some other choice? Figure 3 illustrates the importance of this choice. For the CS log-linearization approach, we graph the stock price elasticities with respect to expected dividends (left) and expected returns (right) for two different choices of ρ. We compare elasticities from calculating ρ based on the value of D/P at the start of 2020 (D/P=1.5%,

$\rho = 0.985$) or based on the historical average dividend-price ratio for the 1871M1-2021M7 sample (D/P=3.9%, $\rho = 0.962$). Using the D/P at the start of 2020 leads to a lower stock price elasticity with respect to expected dividends for early maturities, while the stock price elasticity with respect to expected returns is substantially more negative (more so for longer maturities) when using this value of ρ. For example, the elasticity with respect to the expected return in year 20 is -0.74 when basing ρ on the January 2, 2020, D/P ratio but -0.46 when ρ is based on the historical average D/P ratio. Rather than having to make a judgement on which ρ to use, our approach uses the observable dividend strip weights out to year 10 along with the observable G_t^L/R_t^L ratio to calculate dividend strip weights past year $t+10$.

One could consider a version of the CS log-linearization in which $\ln\left(1+\exp\left(d_{t+k}-p_{t+k}\right)\right)$ was log-linearized around $E_t\left(d_{t+k}-p_{t+k}\right)$. Then

$$\frac{dp_t}{dE_t r_{t+k}} = -\rho_{t+1}\rho_{t+2}...\rho_{t+k-1}$$

with

$$\rho_{t+1} = \frac{1}{1+\exp\left(E_t\left(d_{t+1}-p_{t+1}\right)\right)},\ \rho_{t+2} = \frac{1}{1+\exp\left(E_t\left(d_{t+2}-p_{t+2}\right)\right)} \text{ etc.}$$

This would be more accurate than the standard CS log-linearization. However, because $d_{t+1}-p_{t+1}$ in ρ_{t+1} is in logs, d_{t+1} does not map directly dividend futures. Furthermore, p_{t+1} is a future price so implementing $E_t\left(d_{t+1}-p_{t+1}\right)$ would require assumptions about price expectations (similarly for ρ_{t+2} etc.).[6][7][8]

In sum, our decomposition makes it clear how to use dividend futures to calculate stock price elasticities rather than having to make assumptions about the appropriate value of ρ to calculate stock price elasticities based on the CS log-linearization.

[6] A disadvantage of the CS approach as rewritten above is that it leads to variance terms which do not map directly to observables. One could work directly with equation (20) to avoid that. One would need a lower bound for expected log returns, as opposed to expected returns, but that is possible based on Gao and Martin (2021).

[7] Gao and Martin (2021) make a related point, showing that the Campbell-Shiller log-linearization can be inaccurate when the log price-dividend ratio is far from its historical average. This issue may be particularly relevant for the year 2020 given the COVID recession. They propose a different log-approach and use it to understand the expected growth rate investors must have to be happy to hold the market at a given point in time.

[8] Gonçalves (2021) uses a VAR approach to allow ρ to vary over time in the CS log-linearization. The log dividend price ratio dp_t is included as an element of the state vector and therefore conditional expectations of this variable can be extracted from the VAR model estimation. Our approach does not require assumptions about the D/P process or the estimation of a VAR.

IV. Real yield curve news

A. Real yield curve data

As a baseline approach we propose calculating real interest rates as the nominal interest rate on interest rate swaps (the fixed rate in a fixed-for-floating interest rate swap) minus the rate on inflation swaps. Both series are available out to 40 years, with data obtained from Bloomberg. We convert the term-structure of real zero-coupon yields into annual forward real risk-free rates enabling us to rely on Result 1 in which the expected return for a given future year is changed. For robustness, we consider real interest rates calculated from nominal Treasury rates minus inflation swaps, with nominal Treasury zero-coupon yields obtained from the Federal Reserve.[9] We also consider real interest rates from TIPS, with data from FRED. Both alternatives are available to the 30-year maturity.

Observability of real interest rates by either of the three methods is a relatively recent phenomenon. Inflation swap data is available from Bloomberg beginning in July 2004 and TIPS yields from FRED start in 2003.

B. Real yield curve news in 2020

Figure 4, top graph, shows the evolution of the 10-year and 30-year real rates estimated from interest rate swaps and inflation swaps. Real yields fall dramatically over 2020, with a 114 bps decline in the 10-year real rate and an 80 bps decline in the 30-year real rate. The bottom graphs in Figure 4 illustrates the role of interest rate swaps (the nominal interest rate) and inflation swaps (capturing expected inflation) separately. Both fell dramatically in at the onset of the crisis, with a larger drop in the nominal rate up to March 9 driving the decline in real rates up to this point. After a short-lived spike in nominal yields from March 9 to 18, nominal yields fall and do not recover fully over the year. By contrast, inflation swaps start increasing in late March, resulting in falling real rates up to August. As a side note, it is interesting to observe that financial markets initially thought the COVID crisis would reduce inflation while in fact it turned out to be inflationary (perhaps due to stronger than expected fiscal and monetary stimulus).

The top left panel of Figure 5 shows how changes in real yields affected the stock market over the year via real yield curve news based on our decomposition. The figure is based on our baseline

[9] https://www.federalreserve.gov/data/nominal-yield-curve.htm

approach of using changes in real yields out to 40 years maturity based on interest rate swaps and inflation swaps. The figure shows that falling real yields had a very important role in the rise in the stock market. In fact, yield curve news, taken in isolation, explains an 20.9% increase in the stock market (that increased +22.1% overall). Given that the stock market is such a long-duration asset, the large fall in long-term real yields had a big impact on its price level.[10]

Figure 6 considers robustness to using the two alternative approached to calculate real rates. The decline in the 30 year real rate is around 80bps whether we compute real yields as the difference between interest rate swaps and inflation swaps, the difference between Treasury nominal yields and inflation swaps, or from inflation indexed Treasuries (TIPS). The real yields are about 50bps lower in terms of levels when measured based on interest rates swaps and inflation swaps. Klingler and Sundaresan (2019) argue that low rates in interest rate swaps is due to certain pension funds preferring to get interest rate exposure off-balance sheet. However, the level difference is not important for our results given that it is the changes (and not levels) of these yields that generate the real yield curve news component of stock return. Indeed, the right panel of Figure 6 shows that yield curve news for the year is broadly similar whether we use interest rate swaps or Treasury yields for nominal yield component of our real yields (and using real yields out to the 30-year maturity).

In the above figures, it is noticeable that the decline in real yields is interrupted by a sharp spike in real long yields from March 9 to March 18. The spike is particularly dramatic for real yields based on nominal Treasuries or TIPS. Vissing-Jørgensen (2021) and He, Nagel, and Song (2022) study this spike which led to Federal Reserve purchases of over $1T of Treasuries in 2020Q1 in order to stabilize Treasury markets. The spike is associated with heavy selling by bond mutual funds, foreign central banks and hedge funds and reverses as the Federal Reserve increases its daily purchases sharply starting on March 19. Although there was price pressure pass through into interest rate swaps, the spike in these were less dramatic (perhaps due to derivatives absorbing less balance sheet space than holding Treasuries outright). It is possible that the March spike in real yields is disconnected from the stock market in the sense that stock market investors viewed it driven by urgent liquidity needs that may have been perceived as less relevant for the fundamental value of stocks. If so, our real yield curve news component based on interest rate swaps will be the most accurate but we will nonetheless overestimate the negative return effect of the spike on

[10] As of the start of 2020, the duration of the stock market was 46 years based on the dividend strip weights in Figure 2

realized stock returns in March 2020. This issue will not affect our decomposition for the full year, nor our assessment of the role of the risk premium for the crash, nor our estimate of the role of the real rates outside of the March period.

V. Equity risk premium news

A. Data for the equity risk premium: The Martin lower bound

For the equity risk premium we use the methodology of Martin (2017) who calculates a lower bound on the risk premium using prices of stock market index options and argues that this lower bound is approximately equal to the true risk premium. Martin's data covers the period 1996-2012. We extend his series to 2020 using data from OptionMetrics. We are able to almost exactly replicate Martin's series over his sample period. The description below provides a recap of Martin's approach and Appendix 2 details our data construction.

Martin (2017) starts from the fact that the time t price of a claim to a cash flow X_T at time T can either be expressed using the stochastic discount factor M_T or using risk-neutral notation

$$\text{Price}_t = E_t\left(M_T X_T\right) = \frac{1}{R_{f,t}} E_t^*\left(X_T\right)$$

where the expectation E_t^* is defined by

$$E_t^*\left(X_T\right) = E_t\left(R_{f,t} M_T X_T\right).$$

The return on an investment can similarly be written in terms of the SDF or using risk-neutral notation

$$1 \quad = \quad E_t\left(M_T R_T\right) = \frac{1}{R_{f,t}} E_t\left(R_{f,t} M_T R_T\right) = \frac{1}{R_{f,t}} E_t^*\left(R_T\right).$$

The conditional risk-neutral variance can be expressed as

$$\text{var}_t^* R_T = E_t^* R_T^2 - \left(E_t^* R_T\right)^2 = R_{f,t} E_t\left(M_T R_T^2\right) - R_{f,t}^2$$

The risk premium expressed as a function of the risk-neutral variance is then

$$\begin{aligned} E_t R_T - R_{f,t} &= \left[E_t\left(M_T R_T^2\right) - R_{f,t}\right] - \left[E_t\left(M_T R_T^2\right) - E_t R_T\right] \\ &= \frac{1}{R_{f,t}} \mathrm{var}_t^* R_T - cov_t\left(M_T R_T, R_T\right) \\ &\geq \frac{1}{R_{f,t}} \mathrm{var}_t^* R_T \text{ if } cov_t\left(M_T R_T, R_T\right) \leq 0 \end{aligned}$$

Thus $\frac{1}{R_{f,t}}\mathrm{var}_t^* R_T$ provides a lower bound on $E_t R_T - R_{f,t}$ if $cov_t\left(M_T R_T, R_T\right) \leq 0$, denoted the "negative correlation condition" (NCC).

Martin (2017) shows that the lower bound $\frac{1}{R_{f,t}}\mathrm{var}_t^* R_T$ can be calculated from put and call option prices as follows

$$\frac{1}{R_{f,t}} \mathrm{var}_t^* R_T = \frac{2}{S_t^2}\left[\int_0^{F_{t,T}} \mathrm{put}_{t,T}(K)\, dK + \int_{F_{t,T}}^{\infty} \mathrm{call}_{t,T}(K)\, dK\right]$$

where S_t is the stock price at t, $F_{t,T} = E_t^*\left(S_T\right)$ is the forward stock price, and K denotes the option strike price. On any date, it is therefore possible to extract a lower bound estimate for each available maturity of expiring options. Consistent with Martin (2017), we use linear interpolation to calculate constant maturity lower bounds, which post-2006 allows estimates out to about two years and 6 months.

B. *Equity risk premium news in 2020*

Figure 7 shows our estimated equity risk premia (the lower bounds) over 2020, by maturity. All risk premia shown in the figure are annualized. The top graph shows that 5-day risk premia peaked above 100 percent in March. For the one-year horizon, the risk premium increases from around 3 percent at the start of the year to above 15 percent on March 18. Across horizons, annualized risk premia for longer horizons rise less.

As a supplementary way to describe the term structure of equity risk premia, Figure 8 graphs the cumulative equity risk premium by maturity for the beginning and end of the year as well as for March 18, they day risk premia peaks. Higher annualized risk premia at shorter horizons translate in to a concave cumulative equity risk premium. At the peak of the crisis, investors required a risk premium of 4.1 percent to invest for a 30-day period and a risk premium of 15.7 percent to invest over the next year.

The top right panel of Figure 5 shows how changes in the equity risk premium over the next two years impacted the stock return over 2020. The upward spike in equity risk premia in March generates a negative realized return effect which accounts for minus 14.3 percentage point of the realized return of minus 26 percent up to March 18. The equity risk premium news effect recovers somewhat from the height on crisis, but still ends the year negative, contributing -4 percentage points to the overall 2020 stock return.

In Section VI.D we explore whether the equity risk premium moved past year two and the implication of such changes for our decomposition results.

C. Is the Martin lower bound a good measure of the equity risk premium?

C.1. The tightness of the Martin lower bound

Martin (2017) documents an average lower bound over the 1996-2012 period of about 5%, close to the equity premium estimates obtained by Fama and French (2002) using average realized dividend (or earnings) growth rates as an estimate of ex-ante expected capital gains. Martin also tests whether the lower bound is a good predictor of the realized excess return, with small intercepts. He estimates the relation

$$\frac{1}{T-t}\left(R_T - R_{f,t}\right) = a + b \times \frac{1}{T-t}\frac{1}{R_{f,t}}\mathrm{var}_t^* R_T \tag{28}$$

and cannot reject the null of $b = 1$, $a = 0$ for any horizon from 1 month to 1 year.[11]

We re-estimate (28) over the 1996-2020 sample as shown in Table I. Over this longer sample, we find that β is higher than one for most horizons, though not significantly so for most horizons. The intercept is close to zero and insignificant across all horizons.

Our decomposition relies on *changes* in equity risk premia. What matters for accuracy is thus whether b is close to one, with the intercept playing no role. The b estimates above one imply that the true risk premium change exceeds that of the change in the lower bound. It is possible, however, that realized excess returns exceeded expected returns over this particular time period, more so in times of stress (high values of the risk-neutral variance). Fama and French (2002) argue that realized returns exceeded expected returns even over a sample as long as 1951-2000.

[11] Martin's defines a variable $SVIX^2_{t\to T} = \frac{1}{T-t}\mathrm{var}_t^*\left(\frac{R_T}{R_{f,t}}\right)$ and his regressor is thus expressed as $R_{f,t}SVIX^2_{t\to T}$.

Cieslak, Morse, and Vissing-Jorgensen (2019) argue that over the post-1994 period, unexpectedly accommodating monetary policy has contributed to much of the realized excess return on the stock market. If the unexpected positive component of realized returns is sufficiently correlated with risk-neutral variance, then an estimated b above one may not imply that changes in the lower bound are smaller than the true changes in the equity risk premium for a given horizon.

Given the lack of conclusive empirical evidence on whether $b = 1$ or $b > 1$ it is relevant to ask what theory says about how the *change* in the bound relates to the *change* in the true equity risk premium.

C.2. Theoretical results for the Martin lower bound in changes

Suppose an underlying state variable s_t changes and that s_t is signed such that $\frac{\partial\left[\frac{1}{R_{f,t}}\mathrm{var}_t^* R_T\right]}{\partial s_t} > 0$. Then

$$\begin{aligned}
\frac{\partial\left[E_t R_T - R_{f,t}\right]}{\partial s_t} &= \frac{\partial\left[\frac{1}{R_{f,t}}\mathrm{var}_t^* R_T\right]}{\partial s_t} - \frac{\partial cov_t\left(M_T R_T, R_T\right)}{\partial s_t} \\
&\geq \frac{\partial\left[\frac{1}{R_{f,t}}\mathrm{var}_t^* R_T\right]}{\partial s_t} \text{ iff } \frac{\partial cov_t\left(M_T R_T, R_T\right)}{\partial s_t} \leq 0
\end{aligned}$$

It follows that the change in the lower bound is, on average, equal to the true change in the risk premium if the regression coefficient b in (28) equals one. If instead $b > 1$ that would suggest that the regressor is positively correlated with the omitted variable $-cov_t\left(M_T R_T, R_T\right)$ implying that $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} < 0$ and thus that the true change in the risk premium is larger than the change in the lower bound.

To assess theoretically whether $b > 1$ is likely, assume conditional log-normality as follows:

$$\begin{aligned}
M_T &= e^{-r_{f,t}+\sigma_{M,t}Z_{M,T}-\frac{1}{2}\sigma_{M,t}^2} \\
R_T &= e^{\mu_{R,t}+\sigma_{R,t}Z_{R,T}-\frac{1}{2}\sigma_{R,t}^2}
\end{aligned}$$

where $Z_{M,t}$ and $Z_{R,t}$ are (potentially correlated) standard normal random variables. Martin (2017) shows that in the log-normal case, the NCC holds iff the conditional Sharpe ratio exceeds the conditional standard deviation:

$$cov_t\left(M_T R_T, R_T\right) \leq 0 \text{ iff } e^{r_{f,t}+\sigma^2_{R,t}} \leq e^{\mu_{R,t}} \iff \sigma_{R,t} \leq \frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}$$

The following result states conditions that allow us to relate the true change in the risk premium to the change in the lower bound.

Result 6 (**True change vs. lower bound change in equity risk premium**).
Suppose an underlying state variable s_t changes such that $\frac{\partial\left[\frac{1}{R_{f,t}} var^*_t R_T\right]}{\partial s_t} > 0$ and $\frac{\partial \sigma_{R,t}}{\partial s_t} \geq 0$. The true change in the equity risk premium is at least as large as the change in the lower bound iff $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} \leq 0$. Under log-normality, it is sufficient for $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} \leq 0$ that
(1) The NCC holds: $cov_t\left(M_T R_T, R_T\right) \leq 0 \iff \frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} \geq \sigma_{R,t}$, and
(2) $\frac{\partial}{\partial s_t}\left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}\right] \geq \frac{\partial \sigma_{R,t}}{\partial s_t}$.

Proof: *See appendix.*

In addition to log-normality, assume CRRA utility,

$$M_T = \beta\left(\frac{C_T}{C_t}\right)^{-\gamma} = e^{\ln\beta - \gamma\ln(C_T/C_t)}$$

with $\ln\left(C_T/C_t\right) \sim N(\mu_{c,t}, \sigma^2_{c,t})$ conditional on information known at t. Define the consumption beta relative to the market as $\beta^C_t = \frac{cov_t(\ln R_T, \ln(C_T/C_t))}{\sigma^2_{R,t}}$. The following result emerges.

Result 7 (**True change vs. lower bound change in equity risk premium**).
In the log-normal CRRA case,

$$\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} = \gamma\beta^C_t \sigma_{R,t}. \tag{29}$$

so the NCC holds if $\gamma\beta^C_t \geq 1$. Furthermore,

$$\frac{\partial}{\partial s_t}\left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}\right] = \gamma\beta^C_t \frac{\partial \sigma_{R,t}}{\partial s_t} + \gamma\frac{\partial \beta^C_t}{\partial s_t}\sigma_{R,t}$$

so it is sufficient for

$$\frac{\partial}{\partial s_t}\left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}\right] \geq \frac{\partial \sigma_{R,t}}{\partial s_t}$$

that $\gamma\beta^C_t \geq 1$ and $\frac{\partial \beta^C_t}{\partial s_t} \geq 0$.

Proof: *See appendix.*

Therefore, the same condition that ensures the NCC holds, $\gamma\beta_t^C \geq 1$, also helps to ensure that the true change in the risk premium is larger than the change in the lower bound. Martin (2017) argues that the NCC is very likely to hold in the log-normal case since the Sharpe ratio based on realized returns has substantially exceeded the realized standard deviation. The additional condition, $\frac{\partial \beta_t^C}{\partial s_t} \geq 0$ holds if β_t^C is constant. This will be the case for an investor who is fully invested in the market, since then $\beta_t^C = 1$. It will also (approximately) be the case for an investor who is not fully invested in the market as long as the the investor has a roughly constant portfolio weight in the market and the covariance between the market and non-market risky assets is roughly constant over time.

Overall, theoretical considerations suggest that to the extent that the Martin lower bound is not exact, the most likely direction of any bias is that the true changes in the equity risk premium are larger than the changes in the Martin lower bounds.

C.3. Does the Martin lower bound comove with asset-manager excess return expectations?

One approach to assess the validity of Martin's options-based approach to calculate equity premium estimates is to compare it to equity premium estimates of asset managers. Dahlquist and Ibert (2021) collect a dataset of equity risk premium estimates of asset managers (and some of the largest investment consultants). The data is based on asset managers' capital market assumptions, posted publicly or provided to clients. For the period 2005-2020, their data has 541 observations of equity risk premia from a total of 45 asset managers such as J.P. Morgan, BlackRock, Franklin Templeton and AQR. The most typical horizons in their dataset are 10 years (45% of the data), 7 years (13%), and 5 years (10%). Dahlquist and Ibert (2021) show that asset managers' expectations appear a lot more rational than expectations of households. Asset managers' perceived equity risk premium is high when the P/E ratio is low – consistent with statistical predictive relations – whereas the literature on behavioral finance tends to find that households' perceived equity risk premium is low when the P/E ratio is low.

Figure 9, Panel A, graphs the 1-year Martin measure and the 10-year asset manager equity premia, with one graph for each asset manager that contributes at least five data points for this investment horizon. The equity risk premium estimates of various asset managers are highly correlated with the Martin (2017) lower bound of the equity risk premium that is computed from option prices. In terms of levels, there are substantial cross-asset manager differences in perceived equity premium levels. We construct an overall asset manager equity premium series (for the 10-year horizon) by taking out asset manager fixed effects. We estimate the following relation

$$EP_{m,t}^{10} = \alpha + \Sigma_{t=1}^{T}\beta_t D(\text{date} = t) + \Sigma_{m=1}^{M}\delta_m D(\text{manager} = m) + u_{m,t} \tag{30}$$

where m denotes manager and calculate the predicted value, excluding manager fixed effects, for each date t, $\alpha + \Sigma_{t=1}^{T}\beta_t D(\text{date} = t)$. Figure 9, Panel B, graphs the resulting asset manager 10-year equity premium series along with the 1-year Martin measure. The correlation is high at 0.72.[12] The high correlation lends credence to the Martin measure and also suggests that asset manager equity premium perceptions are consistent with actual asset prices. In Section VI.D, we will exploit the fact that the asset manager equity premia are available for a longer maturity than the Martin series to assess the effect of equity premium news past year 2 for our stock return decomposition for 2020.

VI. Near-term dividends, expected returns, and the return decomposition residual

Turning to the remaining components in the decomposition, we next consider the role of news about dividends out to year 10, of expected returns, and of the decomposition residual for understanding the realized stock market return. We lay out general principles and then provide the results for the application to 2020.

A. Near-term dividends

The relationship between dividend futures and dividend strip prices was discussed in Section II.E with:

$$\frac{F_{n,t}}{\left(1 + y_{n,t}^{\text{nom}}\right)^{n}} = P_t^{(n)} = \frac{E_t\left[D_{t+n}\right]}{R_{t,n}}. \tag{31}$$

A change in a discounted dividend futures price therefore reflects either a change in expected dividends or a change in the cumulative discount rate. Since we have data on futures prices and discount rates, we can rearrange to isolate and estimate expected dividends. Rearrange equation

[12] Regressing the 10-year asset manager series on the 1-year Martin measure results in a regression coefficient of 0.41 with a t-statistic of 9.58.

(31) in terms of expected dividends

$$E_t D_{t+n} = R_{t,n} P_t^{(n)} \tag{32}$$

and extract a lower bound on expected dividends

$$E_t D_{t+n} \geq \hat{R}_{t,n} P_t^{(n)} \tag{33}$$

where $\hat{R}_{t,n} = R_{t,n}^F + \frac{1}{R_{t,n}^F} var_t^* \, R_{t,n}$ is the lower bound of the cumulative discount rate computed from the observed risk-free rate and the observed Martin (2017) lower bound of the equity risk premium.[13] We implement equation (33) using real discount rates $\hat{R}_{t,n}$ in order to estimate real dividends, assuming the lower bound is tight and that the risk premium on each dividend strip is the same as that of the market.[14]

Figure 10 shows the constant maturity expected dividend 1-3, 5 and 8 years ahead over the course of 2020. The left figure shows nominal expected dividends and the right hand side shows real expected dividends. The year-1 expected real dividend fell by 33% from January 2nd to its lowest point on April 3. It ended the year down 7% relative to the start of the year expectation. The moves in 2/3 year dividends were similar but longer term dividends are less dramatic, with the year 8 expected real dividend fairly stable in the early part of the year, followed by a subsequent decline. This pattern is quite similar to the modest and gradual decline in the stock price driven by earnings news documented based on analyst forecasts in Landier and Thesmar (2020).

As the first 10 years of dividends account for less than 20% of total stock value, and since only the first few expected divididends fell dramatically in March 2020, near-term dividend fluctuations can account for little of the crash and recovery in the market. We show this in the left middle panel of Figure 5.

[13]The Martin measure has a maximum maturity of two years. To estimate cumulative discount rates at maturities of 3 years and above, we use actual Martin estimates to two years and then assume forward risk premium beyond this maturity are equal to the time series average of the one year risk premium.

[14]Gormsen et al. (2021) take a different approach to calculate the risk premium for each dividend strip. They express the expected return on dividend strip n as $\hat{R}_{t,n} = R_{t,n}^F + \frac{1}{R_{t,n}^F} cov_t^* \left(R_{t,n}, R_{t,n}^{(n)} \right)$. To estimate the risk-neutral covariance of the n-period return on n-period dividend strip with the n-period return on the market they estimate the risk-neutral standard deviation of the dividend strip return (which is available based on short-dated EuroStoxx dividend future options) and consider alternative assumptions for the correlation of the dividend strip return with the market return. They implement this on EuroStoxx data but dividend future options are not available for the US so we cannot use their approach for the US.

B. *Expected returns*

An exact expected return calculation requires daily estimates of expected returns. However, the one-day equity risk premium is not observed. We therefore approximate the daily expected return by using the one-month expected return on each day (using our 1-month Martin risk premium plus the one-month T-bill yield) and converting this to a per day number. This approach, while not exact, captures time-variation in the daily expected returns within the year.[15]

The middle right panel of Figure 5 shows the cumulative expected return on the market. Over the year, the nominal expected return contributed +6.2% to the overall return. Nearly all of this is expected excess returns, with only 0.25% generated from the risk-free rate. The slope of the expected return increases in March and April, as the spike in short-dated risk premia (documented in Figure 7) increased expected returns significantly.

C. *The decomposition residual*

The bottom left panel of Figure 5 reports the combined effect of our observables described in the above sections. The residual (or unexplained) component of the stock market return is then presented in the bottom right panel. The residual captures the effect of dividend expectations past year 10 and any changes to risk premia past year 2 and real riskless rates past year 30. We therefore call it long-term news. The long-term news component is large in the crisis, contributing about 40 percentage points to the crash and a roughly equal amount to the recovery.

Figure 11 pulls all the results described above together. To summarize our key results, the decline in real riskless rates generate a plus 20.9 percent return component for the stock market for 2020 as a whole. It is central for understanding the market's impressive performance for the year. A spike in equity risk premia out to two years accounts for minus 14.3 percentage points of the realized return of minus 26 percent up to March 18, and for -4 percentage points of realized returns over the full year. Near-term dividends play a minor role, which is not surprising since most of the stock market value comes from later dividends. The expected real return contributed 6 percent to the year's rise. The residual term was very important in the March/April crash and recovery, but not important for the year overall.

[15]The expected return is nominal since we are decomposing the nominal equity return for the year.

D. What drove the residual in 2020?

D.1. Movements in the risk premium past year 2

Figure 12, top, illustrates the time series for the equity premium for year 1 and the forward equity premium for year 2 (both annualized). The forward equity premium for year 2 moves up in March, though much less than the equity premium for year 1. Both remain substantially higher at the end of 2020 compared to their values at the start of 2020.

Given that there is some increase in the risk premium even for year 2, it is likely that risk premia increase to some extent even past this horizon. The asset manager expectations dataset of Dahlquist and Ibert (2021) provides us with longer-dated equity risk premium estimates. We use the 10-year asset manager equity premium series to assess the importance of equity premium movements past year 2. We assume that asset managers agree with the Martin-series in year 1 and 2 and calculate the forward equity premium for year 3 to 10 using this assumption and the 10-year asset manager equity premium series (for simplicity, we assume equal forward equity premia in each of years 3 through 10).

We focus this exercise on quarter-end values because 2/3 of the asset manager data are as of the end of the quarter, making data for other dates less reliable.[16] Figure 13, bottom, shows the 10-year asset manager equity premium, the year 3-10 asset manager forward equity premium, and the equity premia for years 1 and 2 based on the Martin approach. The year 3-10 asset manager forward equity premium is a bit below the 10-year asset manager equity premium (due to the high equity premium for year 1). It spikes substantially (by about 1.7 percentage points) over the first quarter of 2020 but ends the year only 0.2 percentage points higher than at the start of 2020.

Figure 13 shows our return decomposition in quarterly data accounting for equity premium changes out to year 10. This results in substantially larger (more negative) risk premium news in 2020Q1 (top right) and therefore substantially less residual news (bottom right). At the end of March, the cumulative effect of the risk premium news is -18.0% (compared to -6.0% in the baseline estimation). Risk premium news for year 1 to 10 can thus account for almost all of the decline in the stock market in 2020Q1 which amounts to -18.0%. The residual remains negative at the end of 2020Q1 because the real yield curve news is substantially positive for 2020Q1. A negative but modest residual up to the end of 2020Q1 is consistent with the prior cash-flow focused

[16] This is partly due to Dahlquist and Ibert (2021)'s (entirely reasonable) assumption that data are end of the prior month when the asset manager only states a year and month and not an exact date.

literature documenting some but modest cash flow news up to this point.

It is important to recognize that our extension to account for risk premium news out to year 10 does not allow us to assess whether it is likely that there is risk premium news past year 10. The asset manager dataset has little data past year 10.

D.2. Yield curve news beyond 40 years

Figure 14 seeks to determine whether changes to real rates past year 40 are likely to be a large component of the residual. We graph the real (annualized) 10-year forward rates for each of the next 3 decades (and 4th decade when using real rates based on interest rate swap data in the first panel). The real forward rate for the 3rd and 4th decade from now fall over the year, though less than the real forward rates for the first decades. Given the decline in the longest observable real rates, it is possible that real rates changed somewhat even past year 40. Indeed, in the UK, inflation-indexed bonds are traded with 50-year maturity and we find that the real forward rate for years 31-50 falls about 40 bps for the year.

To understand how much real yield curve news past year 40 could have affected the stock market over 2020, define

$$\text{long-term real yield curve news} = \sum_{k=41}^{\infty} -\left[\Sigma_{n=k}^{\infty} w_t^{(n)}\right] \frac{\partial f_{t+k}}{E_t R_{t+k}}$$

Suppose all forward real rates past year 40 moved by the same amount as the forward real rate for year 40. Then

$$\text{long-term real yield curve news} = \frac{\partial f_{t+40}}{E_t R_{t+40}} \sum_{k=41}^{\infty} -\left[\Sigma_{n=k}^{\infty} w_t^{(n)}\right]$$

where the long-term weighting factor is[17]

$$
\begin{aligned}
\sum_{k=41}^{\infty}\left[\Sigma_{n=k}^{\infty} w_t^{(n)}\right] &= -\left(w_t^{(41)} + 2w_t^{(42)} + 3w_t^{(43)} + ...\right) \\
&= -w_t^{(10)}\left(\frac{G^L}{R^L}\right)^{30}\left(\frac{G^L}{R^L} + 2\left(\frac{G^L}{R^L}\right)^2 + 3\left(\frac{G^L}{R^L}\right)^3 + ...\right) \\
&= -w_t^{(10)}\left(\frac{G^L}{R^L}\right)^{31}\frac{1}{\left(1-\frac{G^L}{R^L}\right)^2}.
\end{aligned}
$$

and the ratio of the long-run gross dividend growth rate to the long-run gross expected return, $\frac{G^L}{R^L}$, is defined in equation (14).

Using time series averages for the inputs over the year 2020, $w_t^{(10)} = 0.017$, $\frac{G^L}{R^L} = 0.98$, and $E_t R_{t+40} = 1.022$. This implies long-term real yield curve news of $\partial f_{t+40} \times -22.27$. A drop in real rates of 40 bps for all maturities longer than 40 years would thus increase the stock price by about 9%. This would be countered by any negative stock market effect of increases in the risk premium at very long horizons as well (which could be one plausible reason for lower real rates via higher precautionary savings).

VII. Applications to the financial crisis and longer sample

In this section, we present our decomposition over the Global Financial Crisis. We also present results over the full time series of available discount rate data (from July 2004). As we do not have dividend futures available from Bloomberg prior to 2017, for these periods we use one and two year year dividend strip weights implied by option prices (van Binsbergen, Brandt, and Koijen (2012)) and then use equation (16) to extrapolate weights beyond two years.[18]

[17] The last equality follows from the following property of geometric sums:

$$
\sum_{k=1}^{\infty} kx^k = \frac{x}{(1-x)^2}
$$

[18] Dividend strip values can be estimated from put-call parity for European call options:

$$
P_t^{(n)} = p_{t,n} - c_{t,n} + S_t - Xe^{r^f_{t,n}}
$$

where $p_{t,n}$ and $c_{t,n}$ are the price of a put option and a call option respectively. Both have strike X and maturity $t+n$. van Binsbergen, Brandt, and Koijen (2012) match calls and puts traded in the same minute using intra-day data whereas we have end of day prices via OptionMetrics. This adds noise to our estimates, which we smooth with a centered 15 day moving average. The r-squared between our estimated $w_{t,1}$ and the van Binsbergen, Brandt, and Koijen (2012) online data

Figure 15 (top) shows our decomposition implemented during the Global Financial Crisis, focusing on the second half of 2008 and on 2009. Similar to the COVID crisis, short dated risk-premium estimates spiked at the height of the crisis in October/November 2008. These were a key driver in the market crash, with increases in risk premium for 1yr and 2yr maturities contributing -20% to the stock market return. Also as with the COVID crisis, the financial crisis risk premium spike mean-reverted over the following months, having little remaining impact on the cumulative stock return by the end of 2009.

In terms of real yield curve news, real rates (measured from interest rate swaps and inflation swaps) decreased in late 2008 and early 2009, providing an offsetting effect relative to the increase in risk premiums.

The residual in the financial crisis had an important role to play in both the market crash in fall 2008 and the additional market decline in the first few months of 2009. Interestingly, there is no negative real yield curve news or equity risk premium news in the first few months of 2009 (out to the maturities we can measure), suggesting that negative cash flow news may have been the main driver of the market decline during this time.

Figure 15 (bottom) shows our decomposition implemented over our full sample of available data. The uncompounded return on the SP500 index (including dividends) over the period July 2004 to December 2020 was +189%, of which +58% was the expected return on the market, +17% was real yield curve news, -4% was equity risk premium news and the residual news was 120%.[19] While we have shown short-dated risk-premium news to be important in crisis periods, with sharp spikes in expected excess returns, risk premia have historically been quickly mean-reverting. This means the long-run impact is of risk premium news has been very small. The real yield curve news was particularly important in the final few years over our sample. In the full sample, there is a large role to play for the residual, which incorporates all cashflow news as well as long-term discount rate news.

in the overlapping sample (1996-2009) is 97%.

[19] For a given market return, absolute price changes in the index are larger when the index is at a higher price level relative to when the index is at a lower price level. Plotting compounded returns therefore makes news in recent time-periods (higher index level) seem disproportionately large relative to news in earlier time periods (lower index level).

VIII. Conclusion

The paper contributes to answering a core question in asset pricing: what is the role of discount rate news versus cash flow news. We focus on decomposing daily stock returns (summing up components to understand longer periods) and contribute two main ideas. First, by taking derivatives with respect to expected returns and expected dividends in the present value formula for the stock price, a simple decomposition of the realized return for a given period can be derived. This allows a decomposition into real yield curve news, equity risk premium news and dividend news. Second, to implement the decomposition, a lot can be observed about the real yield curve and the equity risk premium in modern financial markets. We illustrate our approach with an application to the US stock market during the COVID crisis in 2020. Movements in the equity risk premium had a substantial role in the market crash and rebound in March and April, while a fall in real risk-free yields far out the yield curve was a key driver of the stock market ending the year with strong positive returns.

References

Bianchi, F. (2020). The great depression and the great recession: A view from financial markets. *Journal of Monetary Economics 114*, 240–261.

Campbell, J. Y. (1991). A Variance Decomposition for Stock Returns. *The Economic Journal 101*(405), 157–159.

Campbell, J. Y. and R. J. Shiller (1988). The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies 1*(3), 195–228.

Cieslak, A., A. Morse, and A. Vissing-Jorgensen (2019). Stock returns over the fomc cycle. *The Journal of Finance 74*(5), 2201–2248.

Cox, J., D. L. Greenwald, and S. C. Ludvigson (2020). What Explains the COVID-19 Stock Market? *MIT Sloan Working Paper*.

Dahlquist, M. and M. Ibert (2021). How Cyclical Are Stock Market Return Expectations? Evidence from Capital Market Assumptions. *Swedish House of Finance Research Paper No. 21-1* (April).

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

Gao, C. and I. W. Martin (2021). Volatility, valuation ratios, and bubbles: An empirical measure of market sentiment. *Journal of Finance 76*(6).

Gonçalves, A. (2021). What Moves Equity Markets? A Term Structure Decomposition for Stock Returns. *University of North Carolina Working Paper*.

Gormsen, N. J. (2021). Time Variation of the Equity Term Structure. *Journal of Finance 76*(4), 1959–1999.

Gormsen, N. J. and R. S. J. Koijen (2020). Coronavirus: Impact on Stock Prices and Growth Expectations. *The Review of Asset Pricing Studies*, 1–34.

Gormsen, N. J., R. S. J. Koijen, and I. Martin (2021). Implied Dividend Volatility and Expected Growth. *AEA Papers and Proceedings, forthcoming*.

Greenwald, D., M. Lettau, and S. Ludvigson (2019). How the Wealth Was Won: Factor Shares as Market Fundamentals. *NBER Working Paper Series* (510), 1–9.

He, Z., S. Nagel, and Z. Song (2022). Treasury inconvenience yields during the covid-19 crisis. *Journal of Financial Economics 143*(1).

Klingler, S. and S. Sundaresan (2019). An explanation of negative swap spreads: Demand for duration from underfunded pension plans. *Journal of Finance 72*(2), 675–710.

Krishnamurthy, A. and A. Vissing-Jorgensen (2011). The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity* (2).

Landier, A. and D. Thesmar (2020). Earnings expectations during the covid-19 crisis. *Review of Asset Pricing Studies 10*(4).

Martin, I. (2017). What is the expected return on the market? *Quarterly Journal of Economics 132*(1), 367–433.

van Binsbergen, J., M. Brandt, and R. Koijen (2012). On the timing and pricing of dividends. *American Economic Review 102*(4).

van Binsbergen, J. H. (2020). Duration-Based Stock Valuation: Reassessing Stock Market Performance and Volatility. *Wharton Working Paper*.

van Binsbergen, J. H., W. Hueskes, R. S. J. Koijen, and E. B. Vrugt (2013). Equity Yields. *Journal of Financial Economics 110*(3), 503–519.

van Binsbergen, J. H. and R. S. Koijen (2017). The term structure of returns: Facts and theory. *Journal of Financial Economics 124*(1), 1–21.

Vissing-Jørgensen, A. (2021). The Treasury Market in Spring 2020 and the Response of the Federal Reserve. *Journal of Monetary Economics 124*, 19–47.

IX. Figures

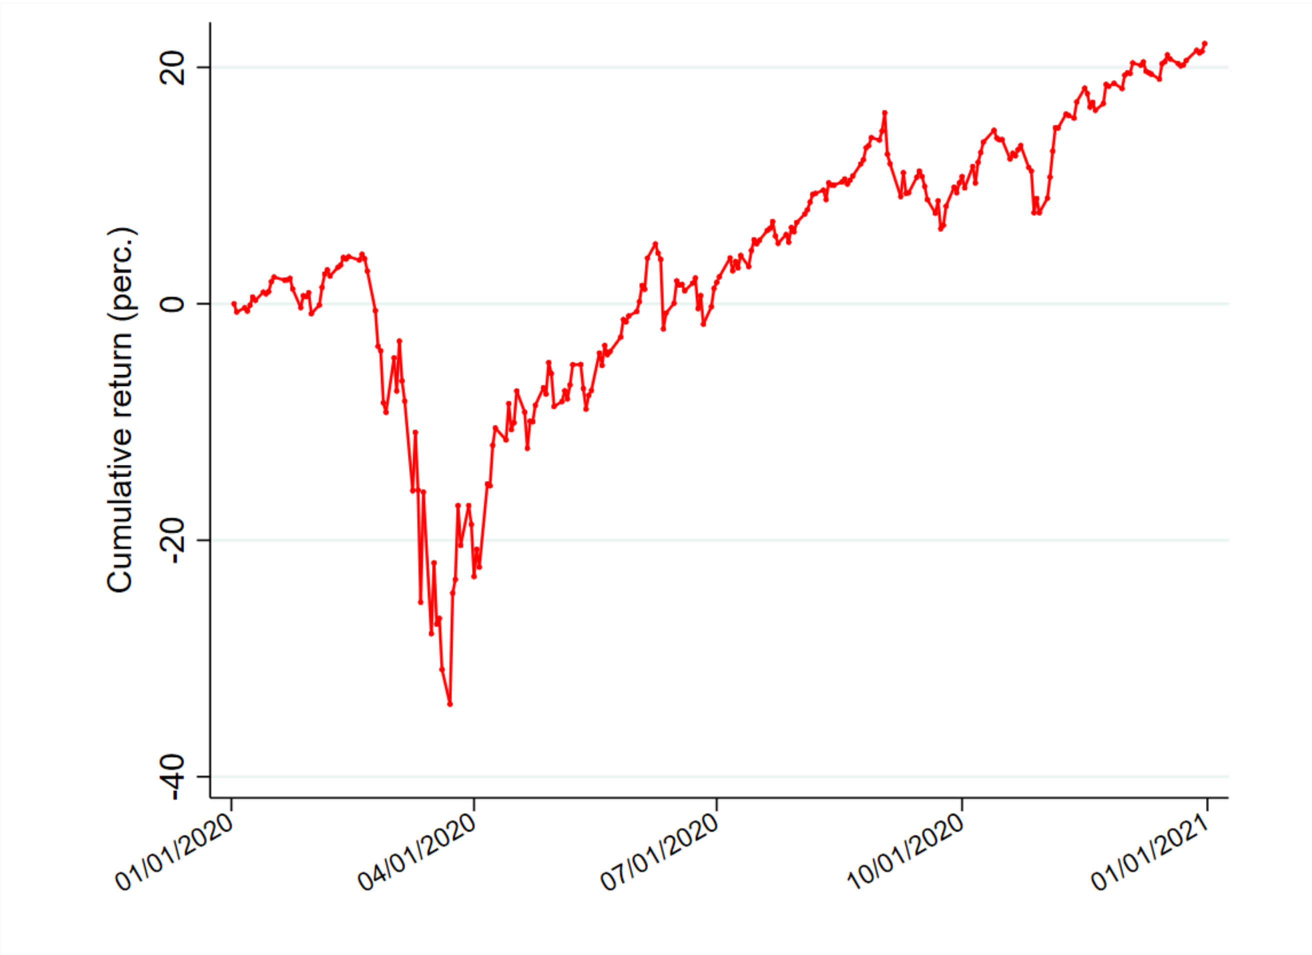

Figure 1. Cumulative return on the S&P500 index, 2020.

This figure shows the cumulative return of the S&P500 index in 2020. It includes price level changes and dividend income. As of March 23rd 2020, the cumulative return year-to-date was negative 33.9%. As of June 2nd 2020, the market had recovered all losses and the cumulative return had turned positive again. The stock market rally continued through the second half of the year and the cumulative return over the full calendar year of 2020 was positive 22.1%.

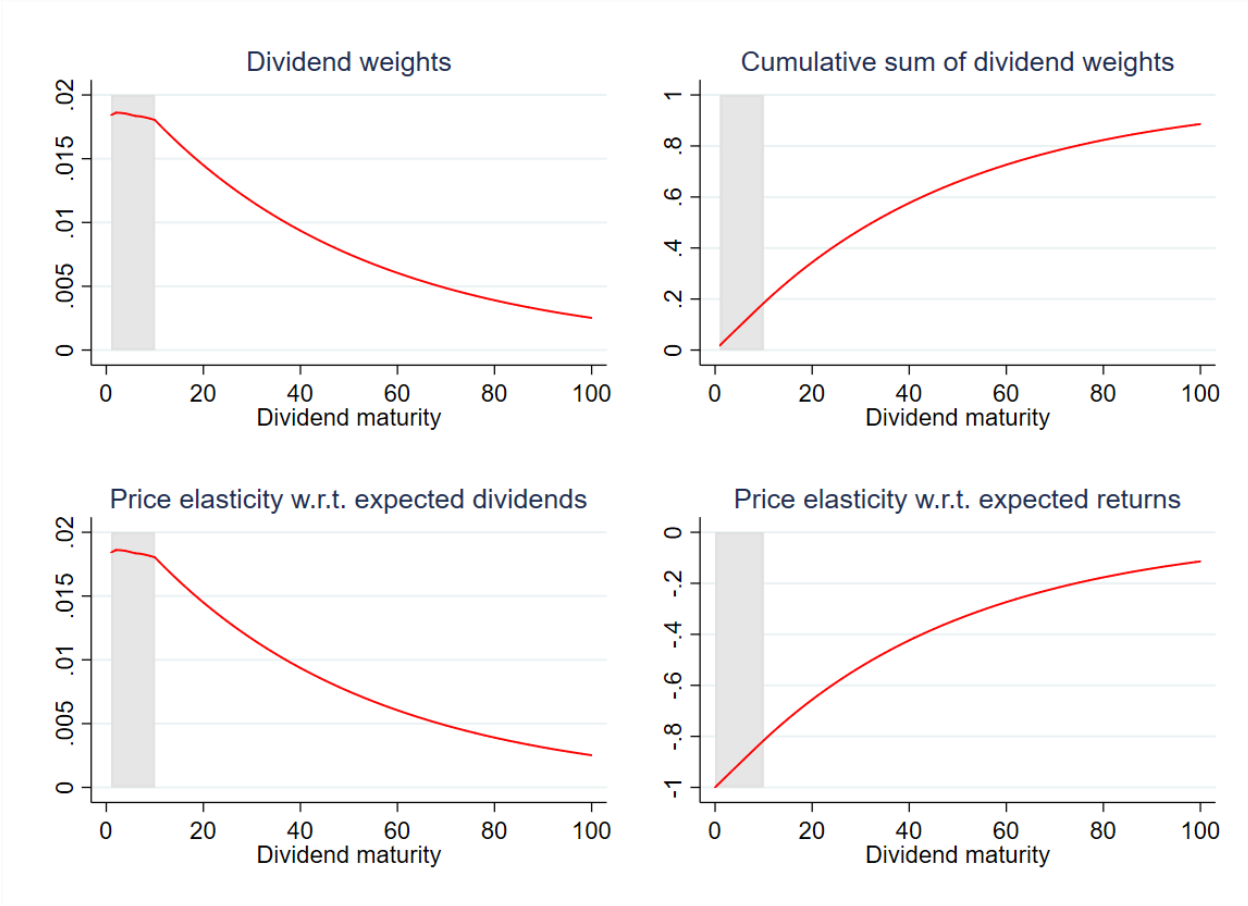

Figure 2. Dividend weights and stock price elasticities.

This figure shows dividend strips weights and how these weights translate into the stock price elasticity with respect to expected dividends and expected returns of various maturities. The top left panel shows the dividend strip weight, $w_t^{(n)}$, for the dividend payment in n years. The weight is the dividend strip's present value as a fraction of the overall stock market valuation. The top right panel shows the cumulative sum of dividend strip weights to maturity n. The bottom left panel shows the stock price elasticity with respect to expected dividend in period $t+k$ as defined in Result 2:

$$\psi_{t,k}^{D} = \frac{\partial P_t / P_t}{\left(\partial E_t D_{t+k}\right) / E_t D_{t+k}} = w_t^{(k)}$$

and the bottom right panel shows the stock price elasticity with respect to expected return in period $t+k$ as defined in Result 1:

$$\psi_{t,k}^{R} = \frac{\partial P_t / P_t}{\left(\partial E_t R_{t+k}\right) / E_t R_{t+k}} = -\sum_{n=k}^{\infty} w_t^{(n)} = -\left(1 - \sum_{n=1}^{k} w_t^{(n-1)}\right).$$

Weights are computed from traded dividend futures prices, which are available to a maturity of 10 years (the shaded region of the figures). From this maturity, dividend weights are estimated using dividend futures prices to year 10 and assuming a Gordon growth model beyond year 10 (see Result 4). All panels use dividend futures prices as of January 2nd 2020.

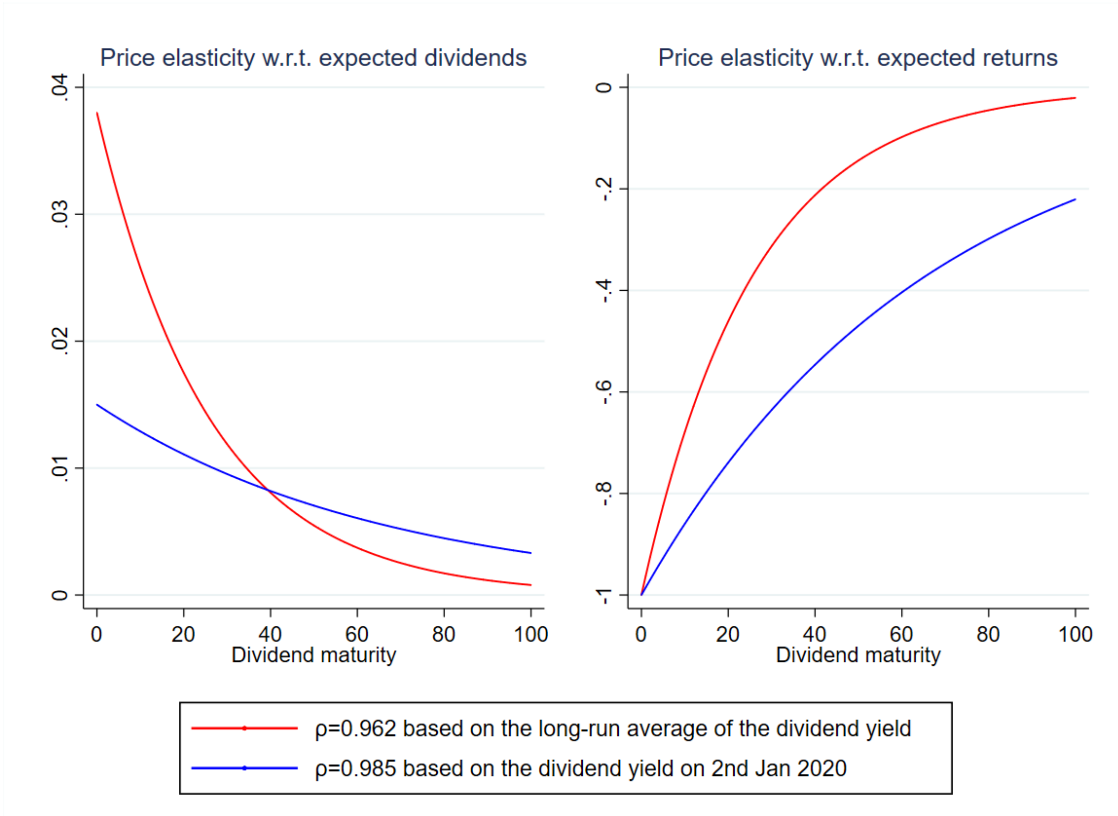

Figure 3. Stock price elasticities using the Campbell-Shiller log-linearization. This figure shows the stock price elasticity with respect to expected dividends and expected returns using the Campbell-Shiller log-linearization approximation. The left panel shows the stock price elasticity with respect to the expected dividend with maturity k, which from Result 2^{CS} is:

$$\phi_k^{D,CS} = (1-\rho)\,\rho^{k-1}$$

and the right hand panel shows the stock price elasticity with respect to the expected return in k periods, which from Result 1^{CS} is:

$$\phi_k^{R,CS} = -\rho^{k-1}.$$

In each panel, we plot the elasticity with two different choices of the dividend yield, $\overline{D/P}$, from which to compute the Campbell-Shiller log-linearization parameter $\rho = \frac{1}{1+\overline{D/P}}$. For the first elasticity, we use the long-run average of the dividend yield ($\overline{D/P} = 3.9\%$, $\rho = 0.962$). For the second elasticity, we use the dividend yield at the start of our main sample on January 2nd 2020 ($\overline{D/P} = 1.5\%$, $\rho = 0.985$).

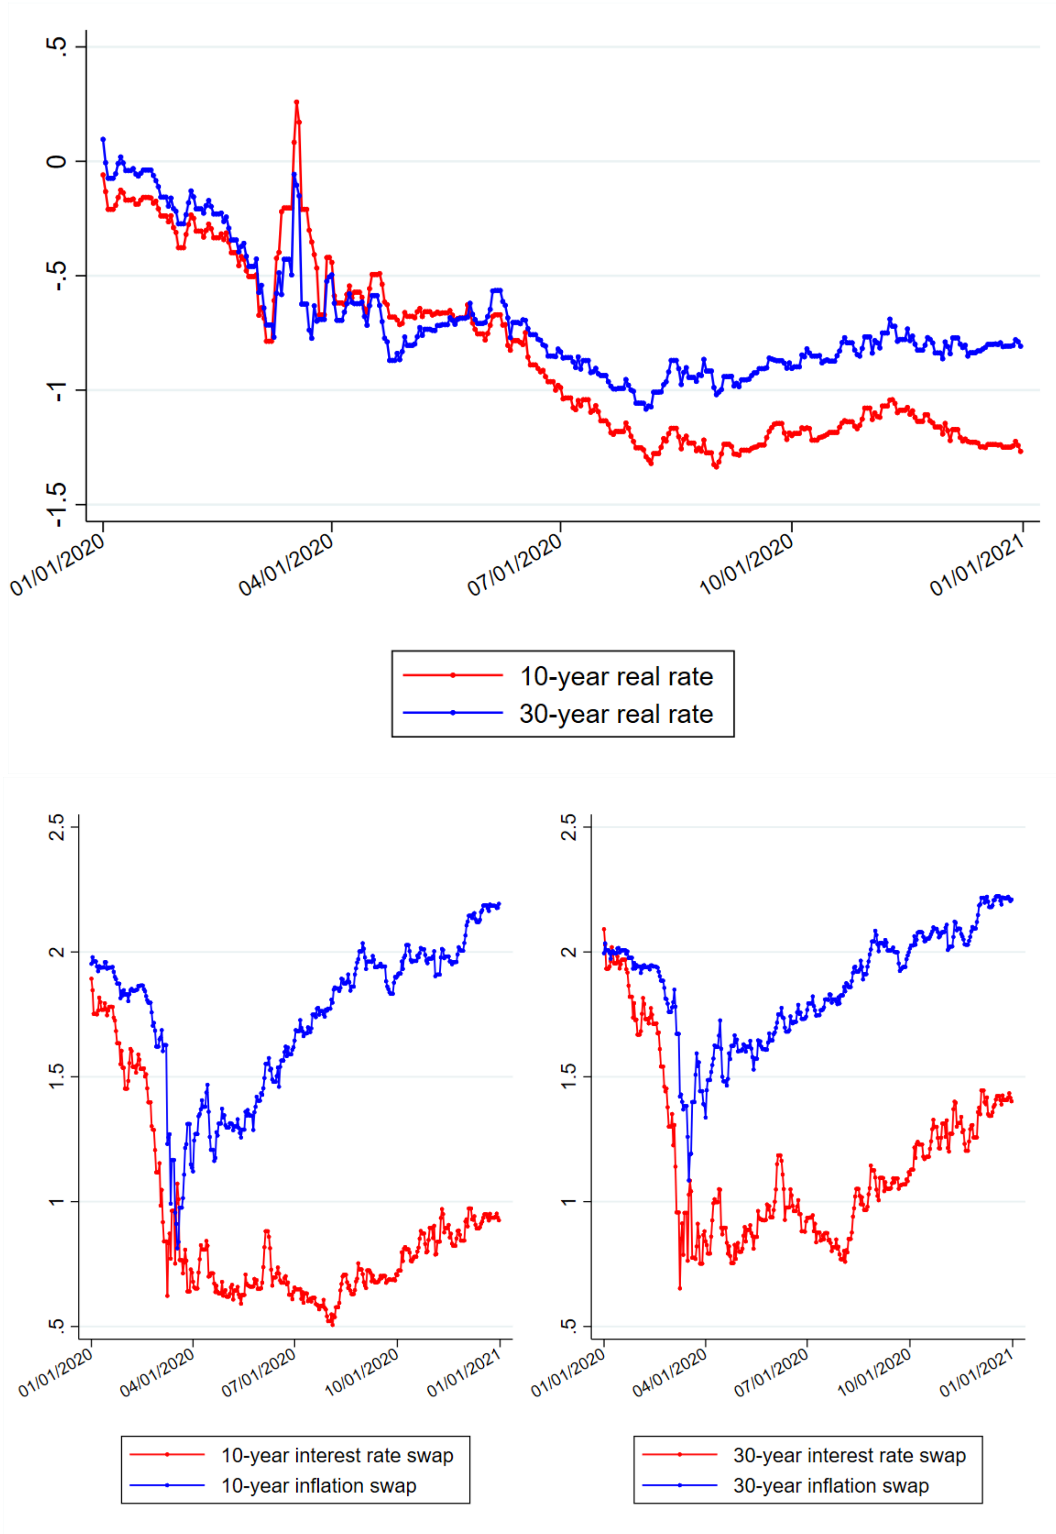

Figure 4. Real risk-free rate and the underlying components.

The top figure shows, for the year 2020, the 10-year and 30-year real interest rates based on (nominal) interest rate swaps and inflation swap rates. The bottom figures show the underlying series separately.

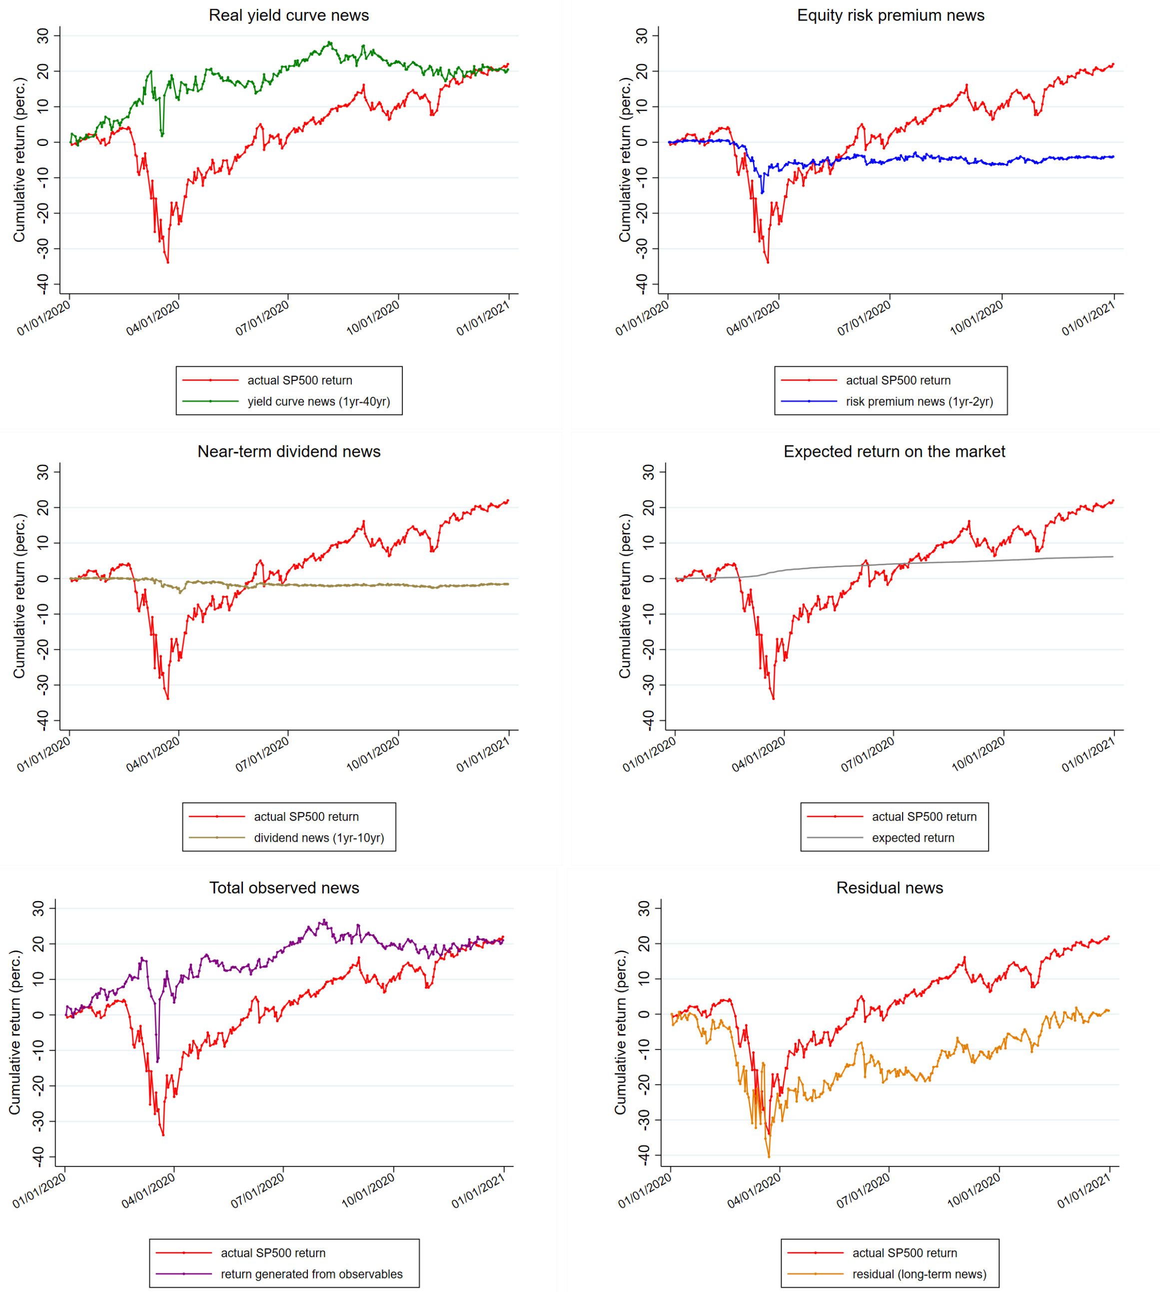

Figure 5. S&P500 return decomposition using observables, 2020

This figure shows the cumulative return of the S&P500 index in 2020, along with the return contribution from yield curve news, equity risk premium news, cashflow news and the realisation of the expected return. All components are extracted from observables in the top 4 panels. The effect of real rate news is estimated from changes in the real risk-free rates (available up to a maturity of 40 years using swaps). The effect of equity risk premium news is estimated from changes in the Martin (2017) lower bound of equity risk premia (available up to a maturity of 2 years using equity option prices). The effect of dividend news is estimated from changes in the price of dividend futures (available up to a maturity of 10 years). The expected return is estimated on a daily basis as the sum of the risk-free rate and equity premium. The sum of the first four panels is the total observable news (shown in the bottom left panel) and the residual (bottom right panel) is the return unexplained by the observable news.

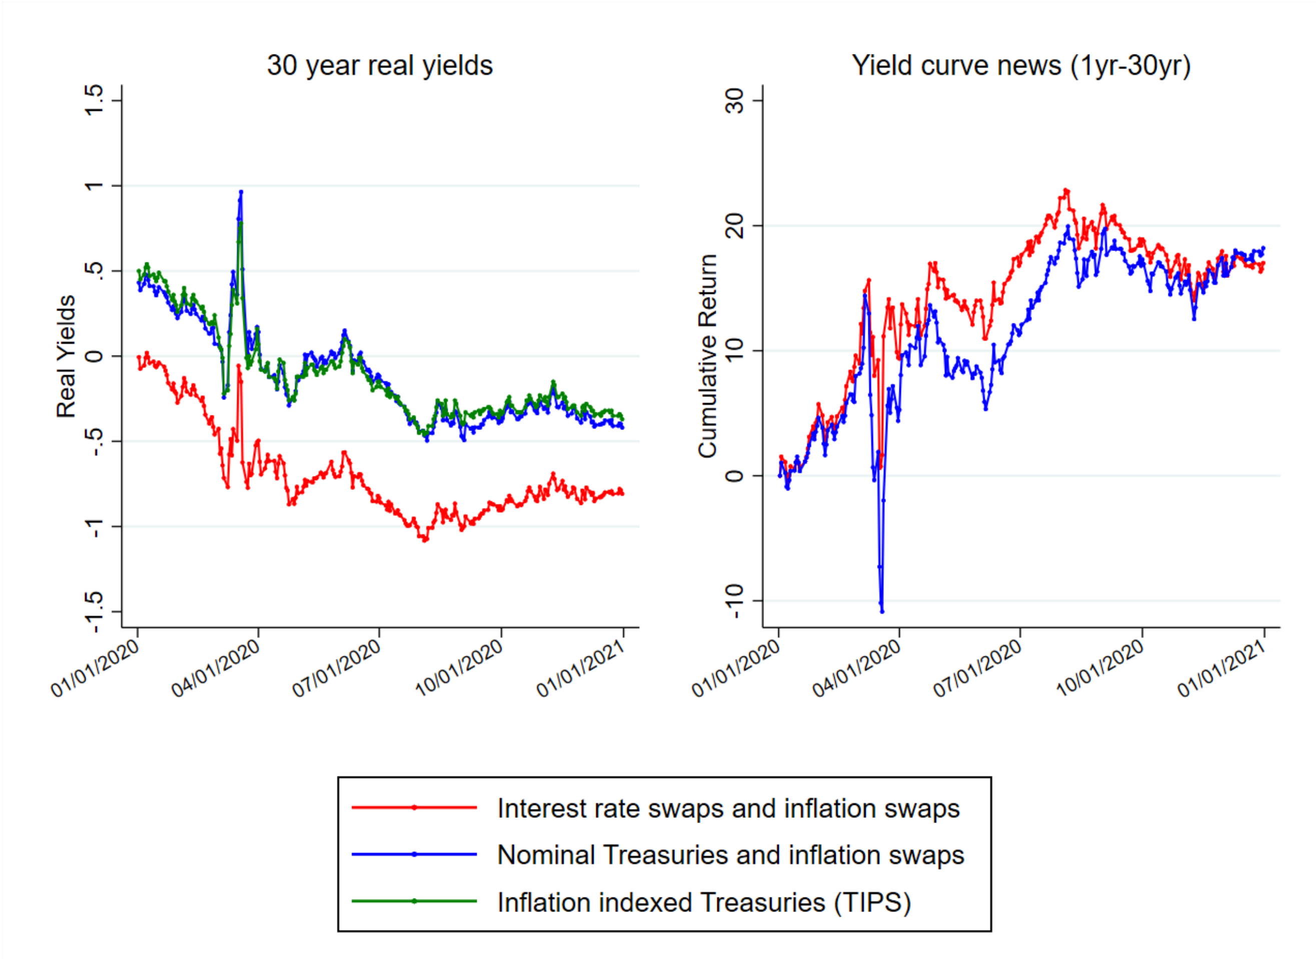

Figure 6. Yield curve news, robustness to different real yield measures

This figure shows that the yield curve news result is robust to different measures of the real yields. The left panel shows the 30 year real yield over 2020 with the real yield calculated using (a) interest rate swaps and inflation swaps, (b) nominal Treasury yields and inflation swaps, and (c) inflation indexed Treasuries (TIPS) yields. The right panel shows the yield curve news using the full yield curve of (a) and (b) out to 30 years. We do not have a full yield curve for (c).

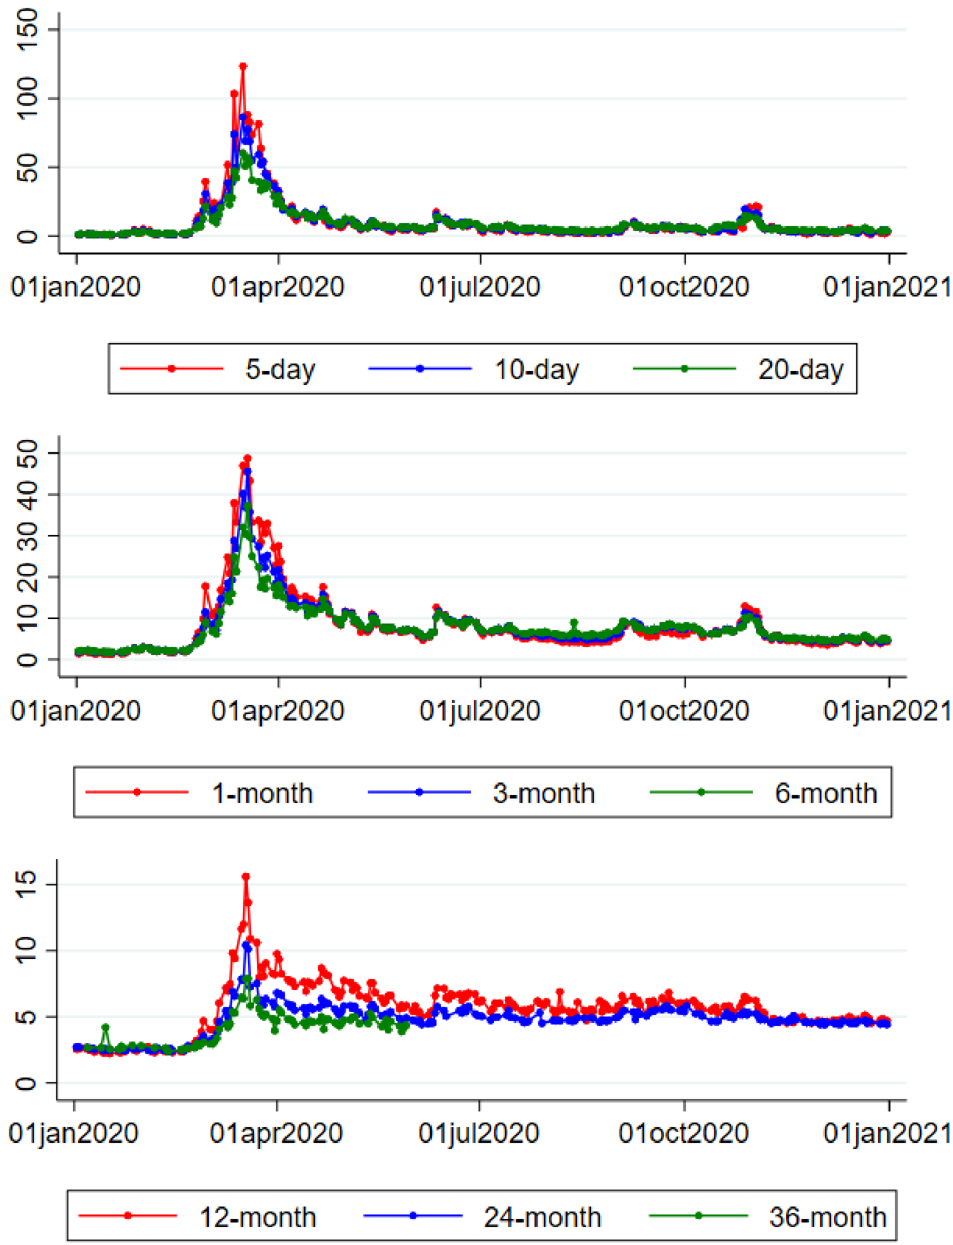

Figure 7. Equity risk premia by maturity (annualized).

This figure shows the annualized Martin (2017) lower bound of the equity risk premium during 2020. It plots the time series of constant maturity risk premium estimates with maturities 5, 10 and 20 days (top figure), 1, 2 and 6 months (middle figure) and 1, 2 and 3 years (bottom figure).

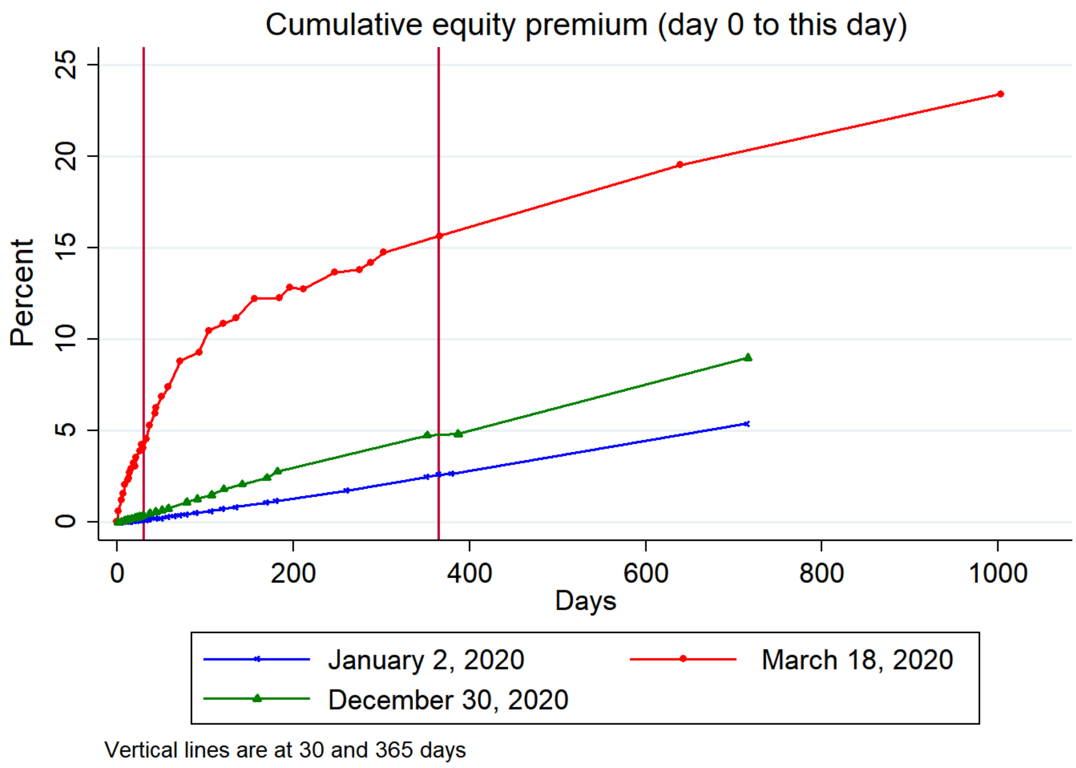

Figure 8. Cumulative equity risk premium.

This figure shows the unanualised Martin (2017) lower bound of the equity risk premium estimates plotted against holding period (expressed in days). It shows the cumulative risk premium curve on three separate dates in 2020: January 2nd (the start of the sample), March 18th (the height of the crisis) and December 30th (the end of the sample).

Panel A.

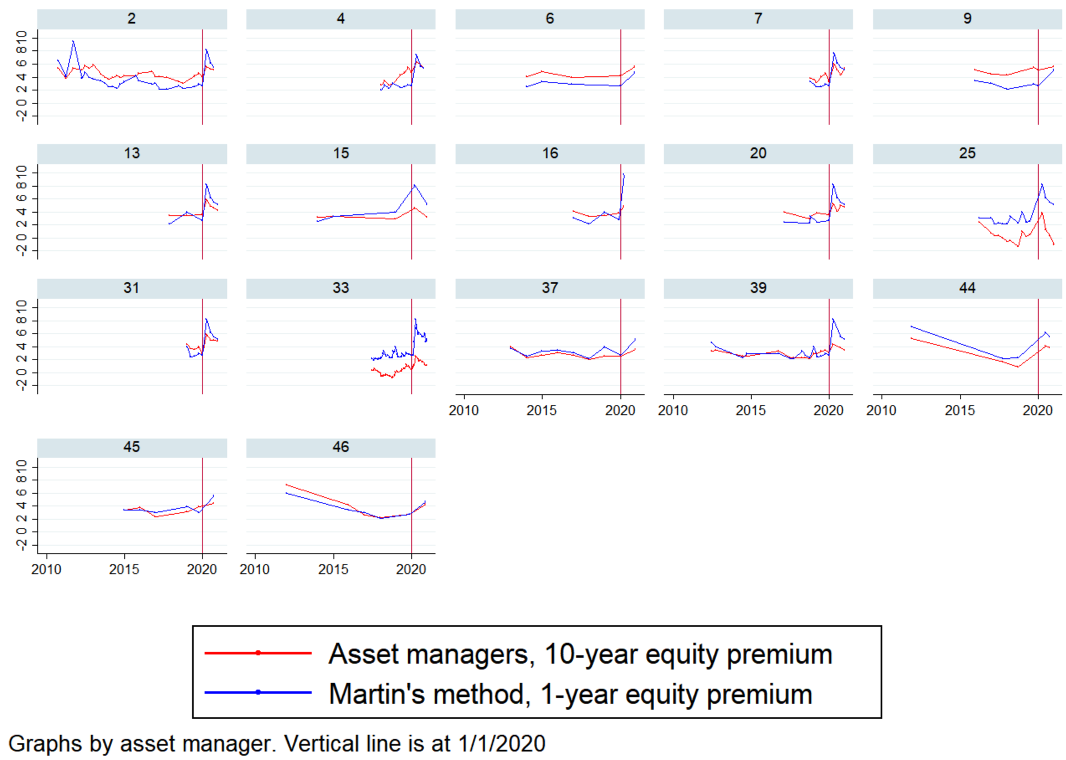

Panel B.

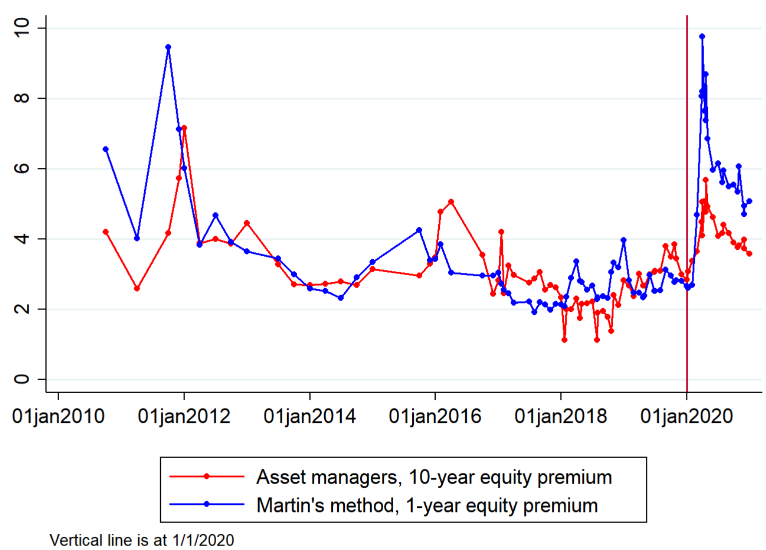

Figure 9.

Asset manager expectations and the option-based equity risk premium estimates

Panel A. This panel shows the 10-year equity risk premium estimates of various asset managers and compares this to the 1-year Martin (2017) lower bound of the equity risk premium that is computed from option prices. Each figure is a time-series plot for one asset manager. Panel B shows a combined series for all asset managers along with the 1-year Martin (2017) lower bound.

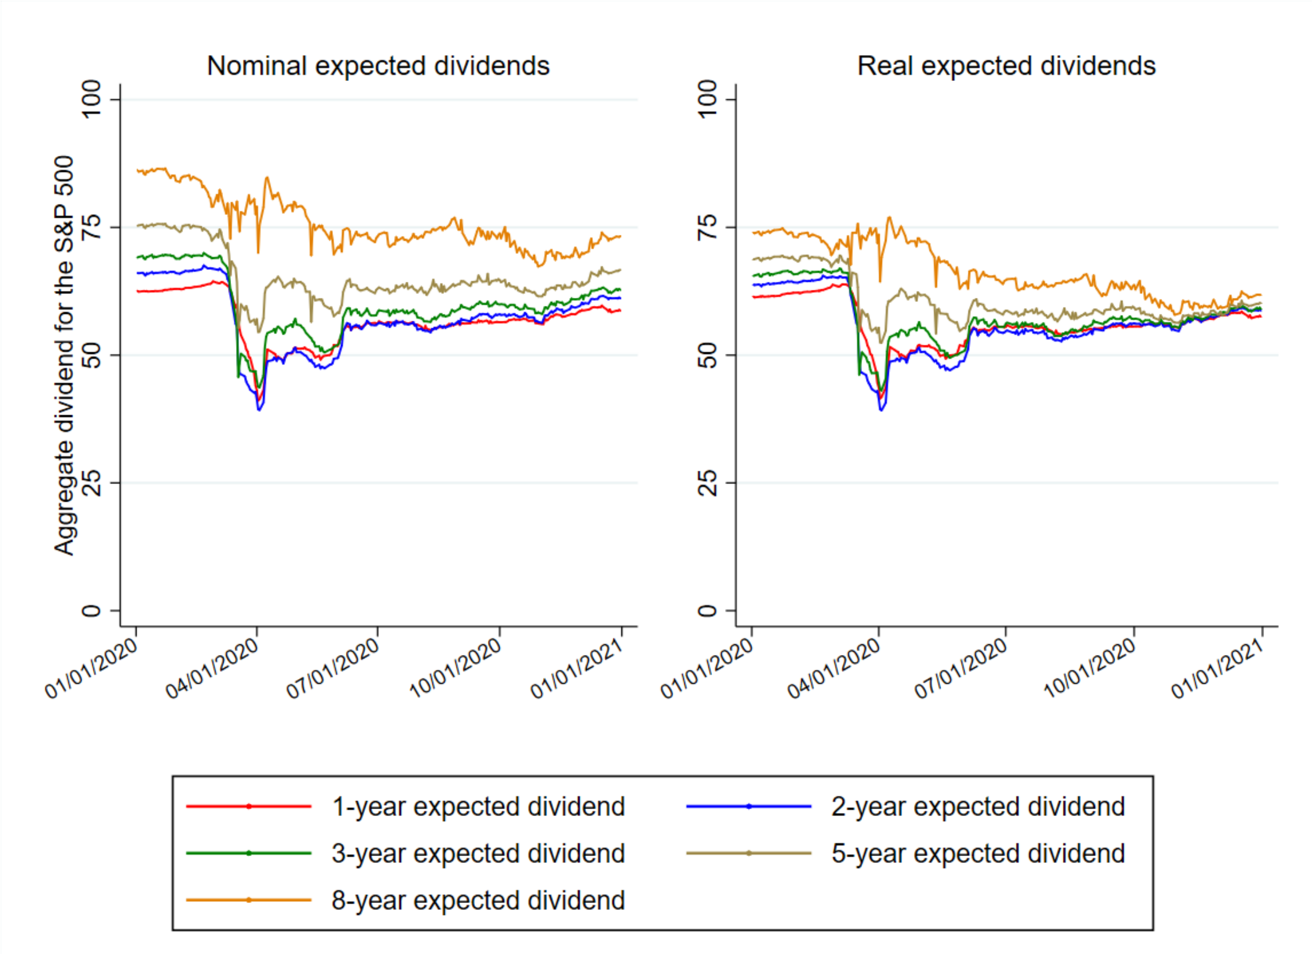

Figure 10. Constant maturity expected dividends

This figure shows dividend expectations extracted from dividend futures. Each figure shows the time-series of constant maturity expected dividends 1-3, 5 and 8 years ahead. The left and right figures show nominal and real expected dividends respectively. Dividends futures prices are adjusted for observable discount rate data to compute an estimation of the expected dividend.

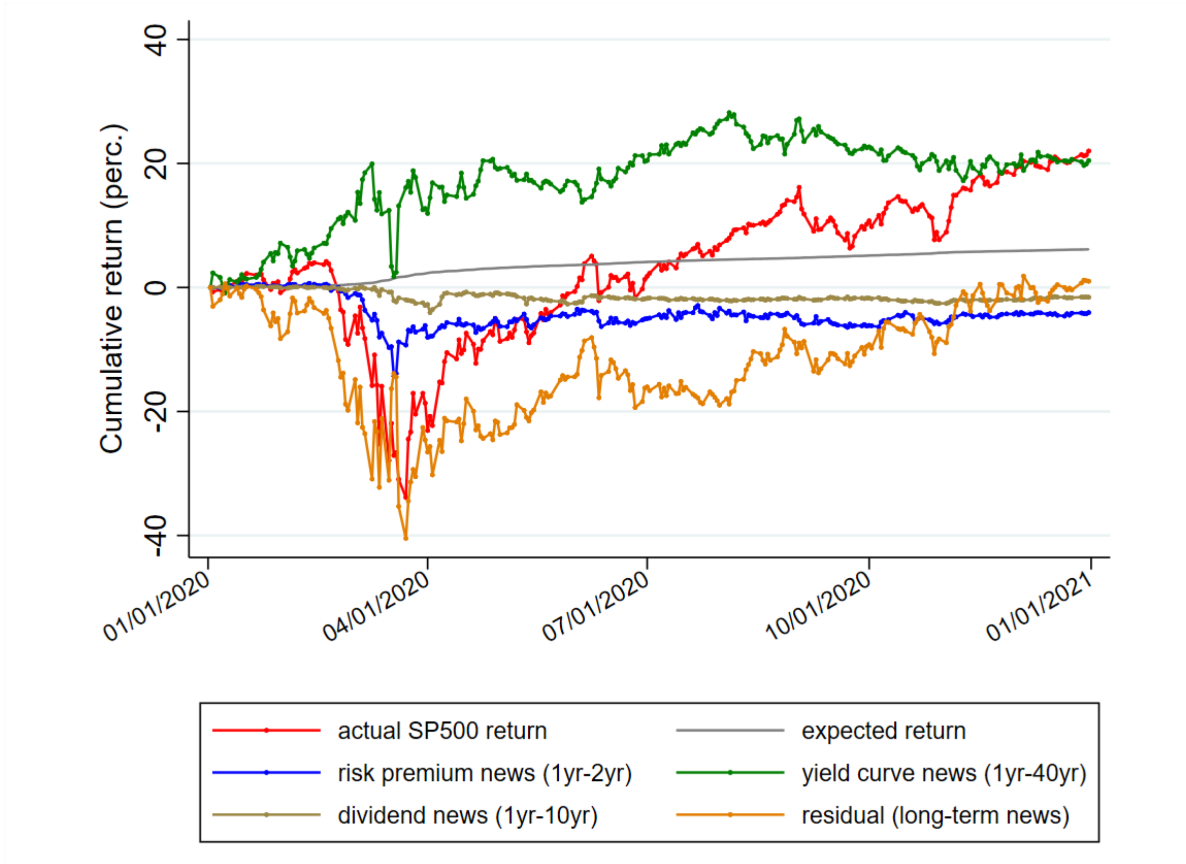

Figure 11. Stock market return from observables and long-term news.

This figure shows the cumulative return of the stock market return along with the implied return from our observables: news about real risk-free rates (1y-40yr), equity risk premium (1yr-2yr), expected dividends (1yr-10yr), and the realisation of the expected return, as well as the residual stock return that is not explained by our observed variables. We call this long-term news.

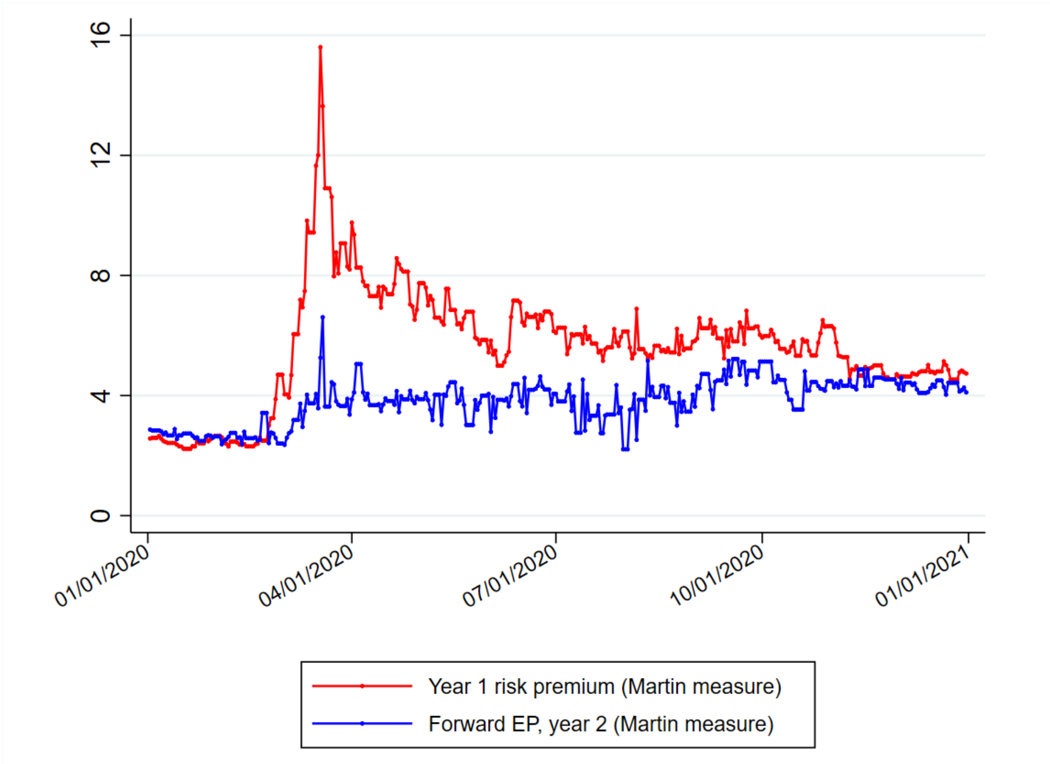

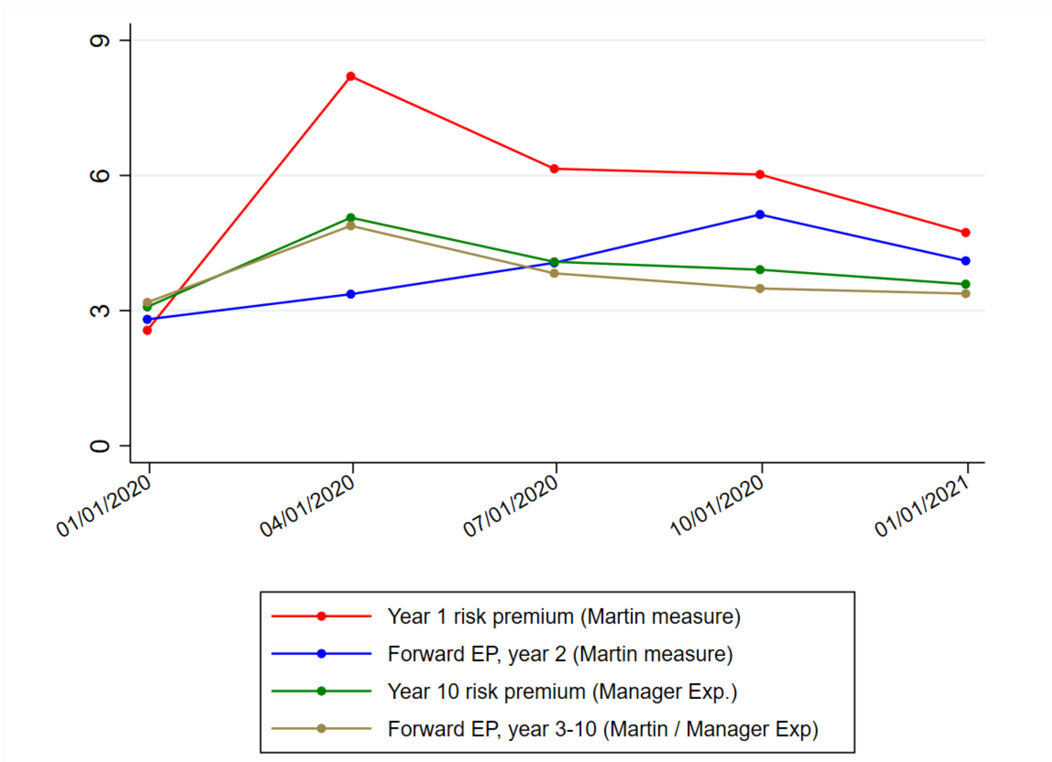

Figure 12. Longer-term equity risk premium estimates.

This figure shows longer-term equity risk premium estimates. The top panel shows one year and one year forward equity risk premiums based from the Martin measure and at a daily frequency. The bottom panel, on a quarterly frequency, shows these variables as well as the 10 year risk premium obtained from manager expectations and the forward risk premium implied through years 3 to 10. The forward risk premium are calculated such that the short-dated Martin measure estimates and the 10-year manager expectations are consistent.

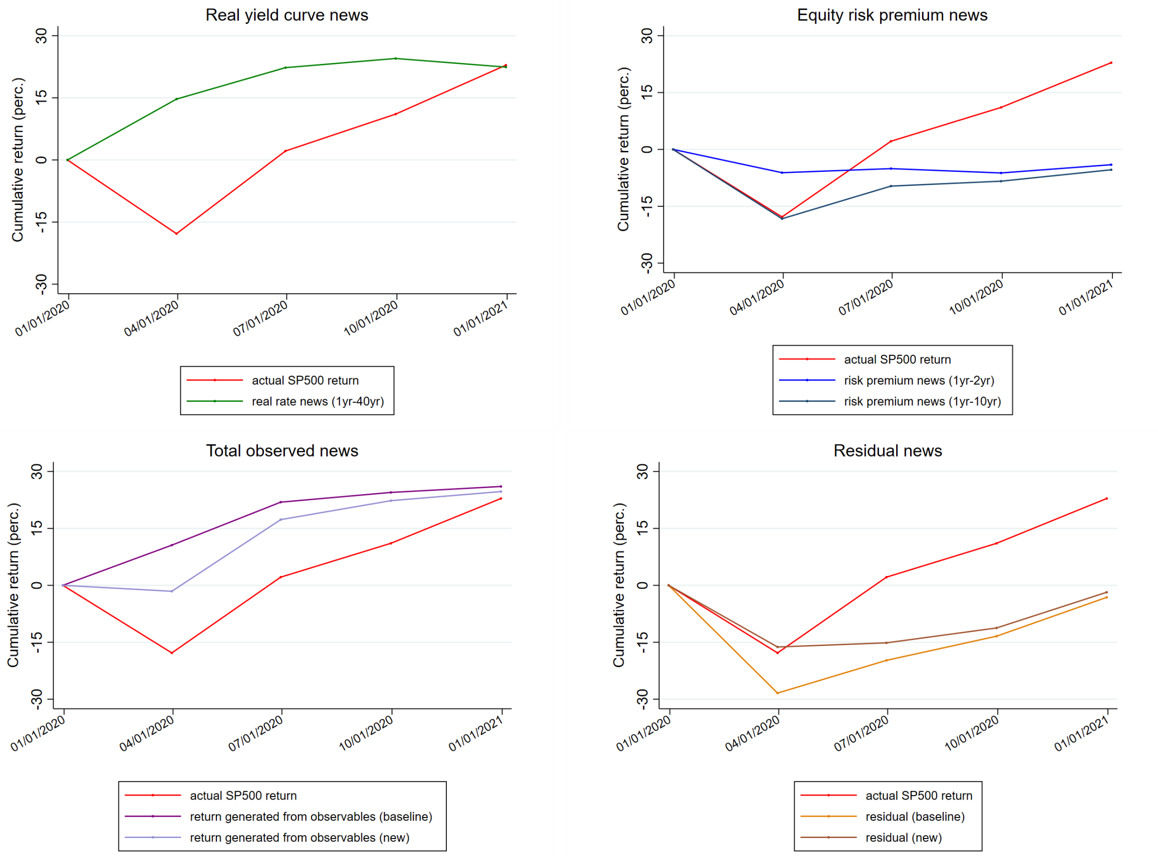

Figure 13. Stock decomposition with longer-term equity risk premium estimates

This figure shows how incorporating longer-term equity risk premium estimates impacts on the decomposition of the cumulative return of the S&P500 index in 2020. The top left shows the return contribution from yield curve news, with the top right panel showing the original equity risk premium news and also the risk premium news using longer-term risk premium maturities as well. The total observable news is shown in the bottom left panel and the residual is shown bottom right panel.

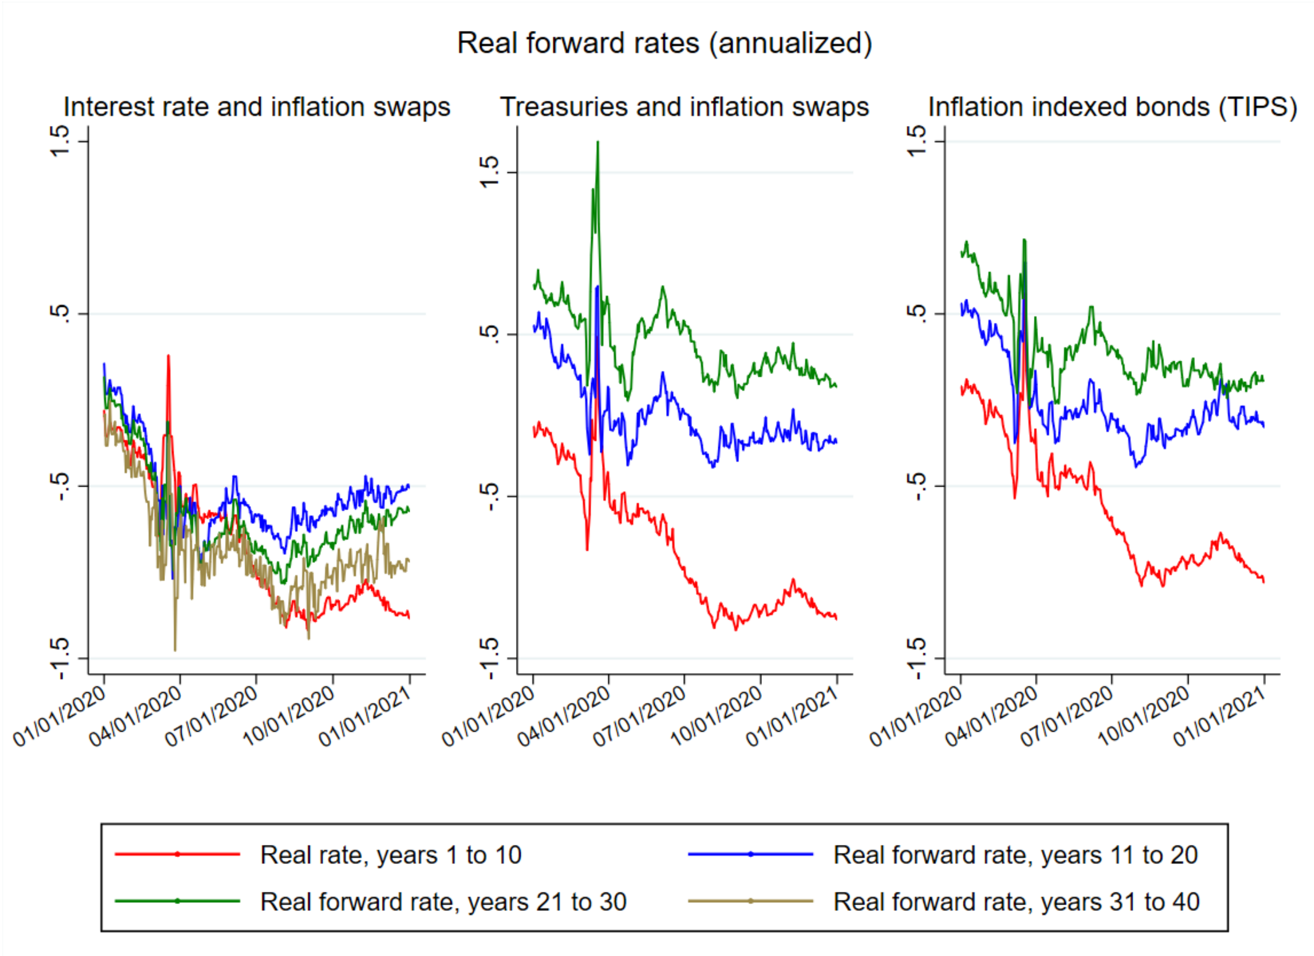

Figure 14. Observed forward real yields

This figure shows real forward rates over 10 year periods as observed from swaps, Treasuries and swaps, and TIPS respectively.

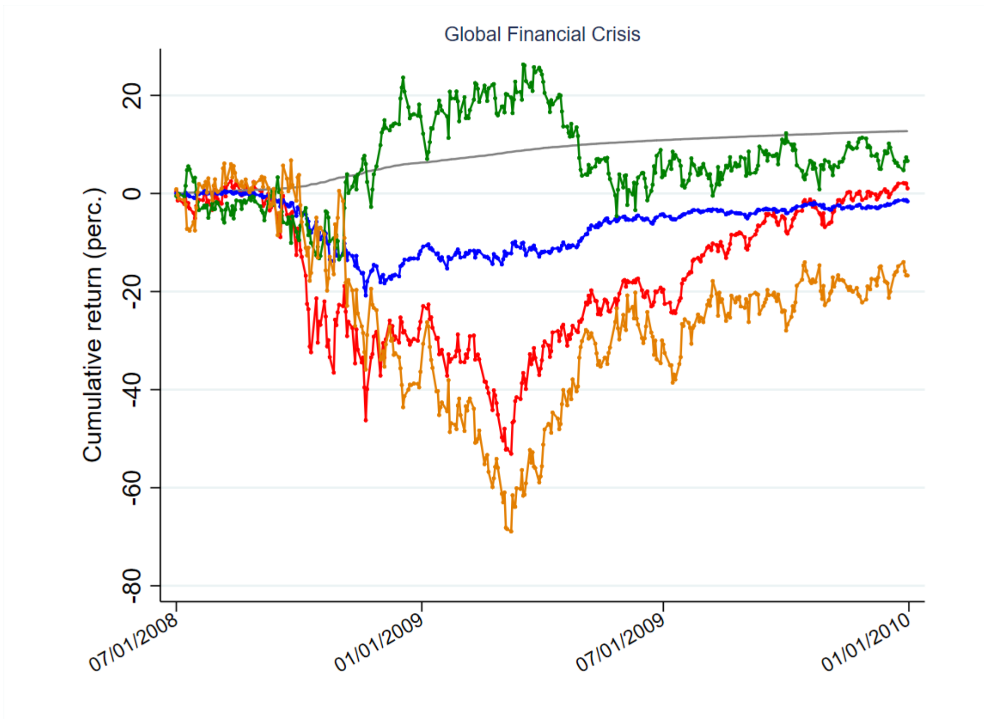

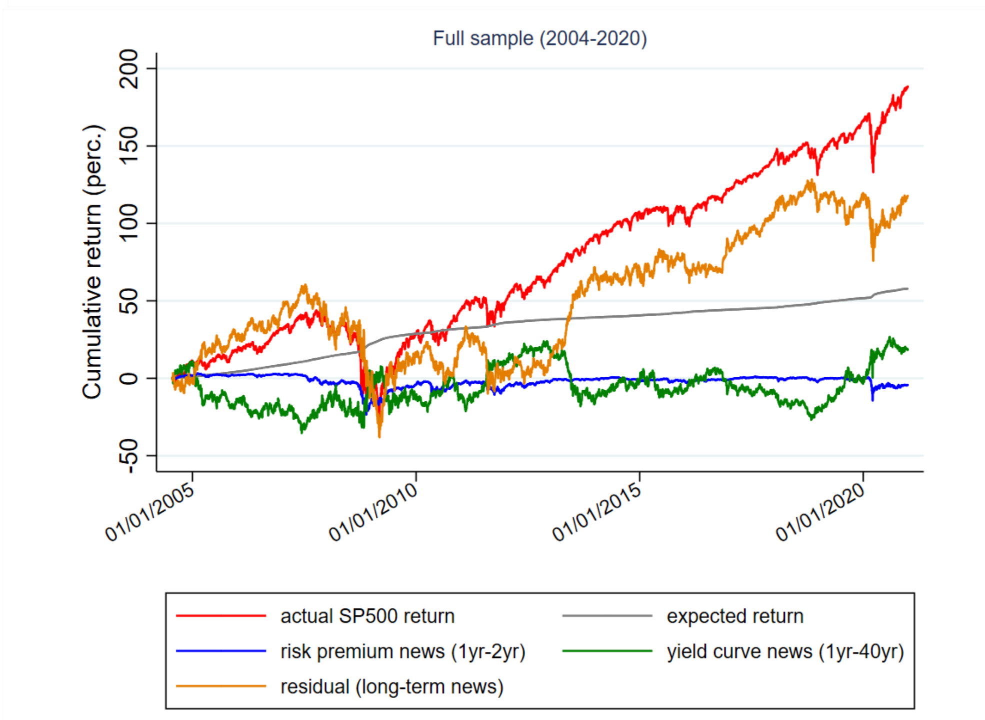

Figure 15. Return decomposition using observables (alternative sample periods)

This figure shows the SP500 return decomposition with observables over other sample periods. The top figure shows the decomposition for the Global Financial Crisis. The bottom figure shows the return decomposition for the full time series of available discount rate data (from July 2004). Over theses periods full dividend futures data is unavailable and therefore the figures don't include a dividends news component. The residual captures all cashflow news as well as long-term discount rate news. Long-term discount rate news includes movements in rates with maturity greater than 2 years for the risk premium component and 40 years for yield curve component.

X. Tables

	Realized Return				
	1 month	2 month	3 month	6 month	1 year
Equity Risk Premium	1.26	1.33	1.33	1.79**	1.09
	(1.01)	(1.04)	(1.27)	(0.88)	(0.88)
Constant	0.01	-0.00	-0.00	-0.02	0.01
	(0.04)	(0.04)	(0.05)	(0.04)	(0.04)
R^2 (percent)	0.96	1.68	2.05	4.64	1.92
Observations	6,254	6,235	6,216	6,151	6,026

Table I. Risk premium estimate as a predictor variable.
This table reports the parameter estimate from the following time series regression: $\frac{1}{T-t}\left(R_T - R_{f,t}\right) = a + b \times \frac{1}{T-t}\frac{1}{R_{f,t}}\text{var}_t^* R_T$ together with Newey-West standard errors with lag selection based on the number of overlapping observations. Columns refer to separate estimations with $T-t = 1$, 2, 3, 6 and 12 months respectively. The sample period is 1996-2020.

XI. Appendix 1. Proofs

A. *Proof of Result 1*

The value of the dividend strip paying D_{t+n} at $t+n$ is

$$P_t^{(n)} = \frac{E_t\left[D_{t+n}\right]}{1+R_{t,n}} \tag{A.1}$$

Using the two assumptions for result 1, it follows that the n-period cumulative discount rate at time t is

$$\begin{aligned}
1+R_{t,n} &= E_t \prod_{k=1}^{n} \left(1+R_{t+k}^{(n-k+1)}\right) && \text{(A.2)}\\
&= \prod_{k=1}^{n} E_t \left(1+R_{t+k}^{(n-k+1)}\right) && \text{(A.3)}\\
&= \prod_{k=1}^{n} b_t^{(n-k+1)} E_t\left(1+R_{t+k}\right). && \text{(A.4)}
\end{aligned}$$

The effect of an instantaneous change to the expected market stock return for year $t+k$ on the value of the dividend strip paying D_{t+n} at $t+n$ can therefore be expressed as

$$\frac{\partial P_t^{(n)}}{\partial E_t R_{t+k}} = \begin{cases} -\frac{1}{E_t[1+R_{t+k}]} P_t^{(n)}, & \text{if } n \geq k, \\ 0, & \text{otherwise.} \end{cases} \tag{A.5}$$

With the overall stock market value $P_t = \sum_{n=1}^{\infty} P_t^{(n)}$, it therefore follows that

$$\frac{\partial P_t/P_t}{\partial E_t R_{t+k}} = \sum_{n=1}^{\infty} \frac{\partial P_t^{(n)}/P_t}{\partial E_t R_{t+k}} = \sum_{n=k}^{\infty} -\frac{1}{E_t\left[1+R_{t+k}\right]} \frac{P_t^{(n)}}{P_t} = -\frac{1}{E_t\left[1+R_{t+k}\right]} \sum_{n=k}^{\infty} w_t^{(n)} \tag{A.6}$$

using the definition $w_t^{(n)} = \frac{P_t^{(n)}}{P_t}$.

B. *Proof of Result 2*

Taking the first derivative of the stock price to an expected dividend we have

$$\begin{aligned}
\frac{\partial P_t}{\partial E_t D_{t+k}} &= \frac{1}{1+R_{t,k}} \\
&= \frac{1}{E_t D_{t+k}} P_t^{(k)}
\end{aligned}$$

and with the definition $w_t^{(k)} = \frac{P_t^{(k)}}{P_t}$ Result 2 follows.

C. Proof of Result 5

With log-linearization around the current D/P ratio (D_t/P_t), and constant expected dividend growth (G) and expected returns (R), ρ becomes

$$\rho = \frac{1}{1 + D/P} = \frac{1}{1 + \frac{R-G}{G}} = \frac{G}{R}$$

Furthermore, with constant expected dividend growth (G) and expected stock return (R)

$$w_t^{(n)} = \frac{P_t^{(n)}}{P_t} = \frac{E_t\left[D_{t+n}\right]/R_{t,n}}{P_t} = \frac{G}{R} w_t^{(n-1)} = \rho w_t^{(n-1)} = \rho^{n-1} w_t^{(1)}$$

Then

$$w_t^{(1)} + \rho w_t^{(1)} + \rho^2 w_t^{(1)} + ... = 1 \Longleftrightarrow w_t^{(1)} = 1 - \rho$$

and

$$w_t^{(n)} = \rho^{n-1} w_t^{(1)} = \rho^{n-1}\left(1 - \rho\right).$$

Therefore,

$$\sum_{n=k}^{\infty} w_t^{(n)} = w_t^{(k)} + w_t^{(k+1)} + ... = \left[\rho^{k-1} + \rho^k + \rho^{k+1} + ...\right]\left(1 - \rho\right) = \rho^{k-1}\left[1 + \rho + \rho^2 + ...\right]\left(1 - \rho\right) = \rho^{k-1}.$$

Result 1 therefore simplifies to

$$\psi_{t,k}^R = -\rho^{k-1}$$

while Result 2 becomes

$$\psi_{t,k}^D = \rho^{k-1}\left(1 - \rho\right).$$

D. Proof of Result 6

$E_t\left(M_T R_T\right) = 1$ implies that

$$\begin{aligned}
\ln\left(E_t\left(M_T R_T\right)\right) &= E_t\left(\ln M_T + \ln R_T\right) + \frac{1}{2} V_t\left(\ln M_T + \ln R_T\right) \\
&= \left(\mu_{R,t} - r_{f,t} - \frac{1}{2}\sigma_{M,t}^2 - \frac{1}{2}\sigma_{R,t}^2\right) + \frac{1}{2}\left(\sigma_{M,t}^2 + \sigma_{R,t}^2 + 2cov_t\left(\ln R_T, \ln M_T\right)\right) \\
&= \mu_{R,t} - r_{f,t} + cov_t\left(\ln R_T, \ln M_T\right) = 0
\end{aligned}$$

and thus

$$\mu_{R,t} - r_{f,t} = -cov_t\left(\ln R_T, \ln M_T\right). \tag{A.7}$$

$E_t\left(M_T R_T\right) = 1$ furthermore implies that

$$cov_t\left(M_T R_T, R_T\right) = E_t\left(M_T R_T^2\right) - E\left(R_T\right)$$

Consider each term on the right hand side separately.

$$\begin{aligned}
\ln E_t\left(M_T R_T^2\right) &= E_t\left(\ln M_T + 2\ln R_T\right) + \frac{1}{2} V_t\left(\ln M_T + 2\ln R_T\right) \\
&= -r_{f,t} - \frac{1}{2}\sigma_{M,t}^2 + 2\left(\mu_{R,t} - \frac{1}{2}\sigma_{R,t}^2\right) + \frac{1}{2}\left(\sigma_{M,t}^2 + 4\sigma_{R,t}^2 - 4\left(\mu_{R,t} - r_{f,t}\right)\right) \\
&= r_{f,t} + \sigma_{R,t}^2
\end{aligned}$$

$$\begin{aligned}
\ln E_t\left(R_T\right) &= E_t\left(\ln R_T\right) + \frac{1}{2} V_t\left(\ln R_T\right) \\
&= \mu_{R,t}
\end{aligned}$$

Combining these two expressions

$$cov_t\left(M_T R_T, R_T\right) = e^{r_{f,t} + \sigma_{R,t}^2} - e^{\mu_{R,t}} \tag{A.8}$$

The derivative with respect to a state variable s_t is

$$\frac{\partial cov_t\left(M_T R_T, R_T\right)}{\partial s_t} = e^{r_{f,t}+\sigma_{R,t}^2}\left[\frac{\partial r_{f,t}}{\partial s_t} + 2\sigma_{R,T}\frac{\partial \sigma_{R,t}}{\partial s_t}\right] - e^{\mu_{R,t}}\left[\frac{\partial \mu_{R,t}}{\partial s_t}\right]$$

If the NCC holds, $cov_t\left(M_T R_T, R_T\right) \leq 0$ and thus $e^{r_{f,t}+\sigma_{R,t}^2} \leq e^{\mu_{R,t}}$. Therefore, it is sufficient for $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} \leq 0$ that $\frac{\partial r_{f,t}}{\partial s_t} + 2\sigma_{R,t}\frac{\partial \sigma_{R,t}}{\partial s_t} \leq \frac{\partial \mu_{R,t}}{\partial s_t}$. Rewrite this sufficient condition as follows

$$\left(\frac{\partial \mu_{R,t}}{\partial s_t} - \frac{\partial r_{f,t}}{\partial s_t}\right)\frac{1}{\sigma_{R,t}} - \frac{\partial \sigma_{R,t}}{\partial s_t} \geq \frac{\partial \sigma_{R,t}}{\partial s_t}$$

Consider now the change in the conditional Sharpe ratio (for log returns):

$$\begin{aligned}
\frac{\partial}{\partial s_t}\left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}\right] &= \frac{1}{\sigma_{R,t}^2}\left[\left(\frac{\partial \mu_{R,t}}{\partial s_t} - \frac{\partial r_{f,t}}{\partial s_t}\right)\sigma_{R,t} - \left(\mu_{R,t} - r_{f,t}\right)\frac{\partial \sigma_{R,t}}{\partial s_t}\right] \\
&= \left(\frac{\partial \mu_{R,t}}{\partial s_t} - \frac{\partial r_{f,t}}{\partial s_t}\right)\frac{1}{\sigma_{R,t}} - \frac{\left(\mu_{R,t} - r_{f,t}\right)}{\sigma_{R,t}^2}\frac{\partial \sigma_{R,t}}{\partial s_t} \\
&\geq \left(\frac{\partial \mu_{R,t}}{\partial s_t} - \frac{\partial r_{f,t}}{\partial s_t}\right)\frac{1}{\sigma_{R,t}} - \frac{\partial \sigma_{R,t}}{\partial s_t}
\end{aligned}$$

where the last line follows from (1) the fact that $\frac{\left(\mu_{R,t} - r_{f,t}\right)}{\sigma_{R,t}^2} \geq 1$ under the NCC and (2) the assumption that $\frac{\partial \sigma_{R,t}}{\partial s_t} \geq 0$. Thus, it is sufficient for $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} \leq 0$ that the change in the conditional Sharpe ratio is at least as large as the change in the conditional standard deviation

$$\frac{\partial}{\partial s_t}\left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}\right] \geq \frac{\partial \sigma_{R,t}}{\partial s_t}.$$

E. *Proof of Result 7*

We can exploit (A.7) to get

$$\begin{aligned}\mu_{R,t} - r_{f,t} &= -cov_t\left(\ln R_T, \ln M_T\right) \\ &= \gamma cov_t\left(\ln R_T, \ln\left(C_T/C_t\right)\right).\end{aligned}$$

This implies (29),

$$\begin{aligned}\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} &= \gamma\frac{cov_t\left(\ln R_T, \ln\left(C_T/C_t\right)\right)}{\sigma_{R,t}^2}\sigma_{R,t} \\ &= \gamma\beta_t^C\sigma_{R,t}\end{aligned}$$

where β_t^C is the (potentially time-varying) beta of $\ln\left(C_T/C_t\right)$ with respect to $\ln R_T$. The rest of Result 7 follows directly from (29).

XII. Appendix 2. Practical issues in implementing the Martin lower bound

For 2020, we use option price data from CBOE to construct a time series of the Martin (2017) lower bound of the equity risk premium. The data includes the trading prices for every traded SP500 Index Option on each day (with intra-day data available), as well as each option's best bid and ask price, strike price, expiry date. The data also reports the underlying SP500 index price at the time of trade. We clean the initial data in several ways. First we delete all trades with a highest bid or ask of zero. Second, we delete trades where the trade price is greater (lower) than the highest ask (bid). Third, we delete all trades where the underlying index price is missing. Finally, from this selection of cleaned trades, we select the latest traded option for each date, expiry date, strike, option type (call/put) combination.

For each date, expiry date, strike combination in the dataset we then estimate the equity risk premium by discretizing the right-hand side of Martin (2017)'s lower bound given by

$$\frac{1}{R_{f,t}}\text{var}_t^* R_T = \frac{2}{S_t^2}\left[\int_0^{F_{t,T}} \text{put}_{t,T}\left(K\right)dK + \int_{F_{t,T}}^{\infty} \text{call}_{t,T}\left(K\right)dK\right].$$

The forward price $F_{t,T}$ is the unique solution K of the equation $\text{call}_{t,T}\left(K\right) = \text{put}_{t,T}$ (K). We estimate the forward price by first interpolating trade prices across strikes for both calls and puts at each date and expiry date combination, and second identifying the intersection of these two interpolated series. We do not use the interpolated prices in discretization of the above equation.

Once equity risk premium estimates have been estimated for each date expiry date combination, we clean the data again. First, we only keep equity risk premium estimates where the number of strikes used in the estimation is greater than 15. Second, we delete estimates when the minimum call or put price is over 40% of the maximum trade price for that date and expiry date combination. These cleaning procedures delete estimates where the discretization is too coarse and where a large tail of options are missing, both of which cause biased estimates.

Finally, we generate constant maturity equity risk premiums by interpolating between those estimated at available expiry dates on any given date.[20] We also extrapolate to extend maturity.

[20] As an alternative to interpolation, we have also estimated the equity risk premium yield curve at each date using cubic

However, to avoid over extrapolation, we limit this extrapolation to a maximum of 150 days greater than the longest maturity option available at that date.

spline methods. The results are very similar.